

Towards a Framework for Teaching Artificial Intelligence in Higher Education

Thesis by

Becky Allen

In Partial Fulfilment of the Requirements for the
Degree of Doctor of Philosophy



School of Computing

Newcastle University

Newcastle Upon Tyne, UK

June 2022

Declaration

All work contained within this thesis represents the original contribution of the author. This research has resulted in a number of publications which are listed below. Most of Chapters 3 and 4 and aspects of Chapter 7 have been published.

List of Publications

Allen, B., McGough, AS., Devlin, M. (2022) Paper presentation at *SIGCSE TS 2022* for: 'Toward a Framework for Teaching Artificial Intelligence to a Higher Education Audience' *ACM Transactions on Computing Education (TOCE)*, <https://sigcse2022.sigcse.org/schedule/>

Crick, T., Davenport, J., Hanna, P., Hayes, A., Irons, A., Miller, K., Prickett, T., Ward, R., Allen, B., Patil, B., Payne, S. (2022) 'Co-Constructing a Community of Practice for Early-Career Computer Science Academics in the UK' (*CEP22*). doi: [10.1145/3498343.3498349](https://doi.org/10.1145/3498343.3498349)

Allen, B., McGough, AS., Devlin, M. (2021) 'Toward a Framework for Teaching Artificial Intelligence to a Higher Education Audience' *ACM Transactions on Computing Education (TOCE)*, 22(2). doi: [10.1145/3485062](https://doi.org/10.1145/3485062)

Allen, B., Heels, L. and Devlin, M. (2021) 'Measuring the Difference Between Student and Staff Perception of Self-Efficacy and Confidence Using Online Tools', in *Proceedings of United Kingdom and Ireland Computing Education Research (UKICER)*. ACM.

Allen, B., Devlin, M., McGough, AS. (2021) 'Using the One Minute Paper to Gain Insight into Potential Threshold Concepts in Artificial Intelligence Courses' *Computing Education Practice (CEP21)*. doi: [10.1145/3437914.3437974](https://doi.org/10.1145/3437914.3437974)

Allen, B. (2019) 'Identifying Best Practices for Teaching Machine Learning,' *ACM-W UK Inspire Conference 2019*. <https://acmukwomen.acm.org/inspire-2019/>

Allen, B. (2019) 'Identifying Best Practices for Teaching Machine Learning,' *CSEDU Doctoral Consortium*.

Abstract

Artificial Intelligence and the domain's sub-disciplines are becoming increasingly prevalent within numerous areas of academia and can now be considered a core area of Computer Science (Shapiro, Fiebrink and Norvig, 2018). As a consequence, the Higher Education (HE) sector are increasing their provision of Machine Learning and Artificial Intelligence courses. However, there is a current lack of research pertaining to the best practice for teaching this complex domain, which relies heavily on both computing and mathematics knowledge.

This thesis outlines a review of the current education provision in AI within higher education, assessed through qualitative techniques encompassing both lecturer and student interactions. Through completion of case studies at varying educational institutions, potential barriers to learning were identified including issues with mathematics anxiety and low confidence in technical skills.

The thesis introduces MetaLearning, a learning resource created as part of this research to serve as an introductory course for Machine Learning. MetaLearning consists of a framework of topics pertinent to an introductory course. This framework was developed from the findings of the review of the current educational provision which identified key topics for inclusion. MetaLearning also incorporates a number of mitigation strategies to assist learners in overcoming some of the identified barriers. Strategies pertain to improving student's metacognition and self-efficacy with the overall aim of learners becoming more self-regulated, therefore equipping them with the tools to persevere when encountering difficulties such as threshold concepts. A review of MetaLearning, outlining both the student and lecturer view of the efficacy of this resource is included.

Finally, an initial framework is outlined for the best practice for teaching AI. This includes issues pertaining to educational background, mathematics anxiety and low self-efficacy. Alongside an initial overview of the potential threshold concepts, guidance to improve student attainment and satisfaction within these courses is also discussed. Although these findings were the outcome of research specific to AI, they have relevance and will generalise to the wider overarching Computer Science domain.

Acknowledgements

I would like to thank my supervisor Dr Marie Devlin for all her support and guidance throughout this project. Thank you for all the opportunities you have generously provided me throughout the research including conferences and workshops. Your support has been vital to my development as a researcher and an academic. I would like to thank my other supervisor Dr Stephen McGough for his advice, assistance and continued support throughout my PhD years. A big thank you to all members of the Educational Practice in Computing group, especially Laura Heels for her regular pep talks and continuous support. I would like to thank all institutions who participated in this research, particularly the academics who generously contributed their time to participate. A special thank you to colleagues at Sunderland University who have supported me through completion of my PhD whilst starting my first academic post, in particular Dr Kate MacFarlane and Professor Sharon McDonald. Also, a big thank you to Jess for all the coffee and cake and chats. Finally, the biggest thank you goes to my family, my Mam, Dad, Kate, Andy and Rosie, I could not have achieved any of this without your unwavering support, thank you for being the best.

Table of Contents

<i>Declaration</i>	<i>i</i>
<i>List of Publications</i>	<i>i</i>
<i>Abstract</i>	<i>ii</i>
<i>Acknowledgements</i>	<i>iii</i>
<i>Table of Tables</i>	<i>x</i>
<i>Chapter 1. Introduction</i>	<i>1</i>
1.1 Research Context	1
1.2 Research Motivation	3
1.3 Research Aim and Research Questions	5
1.4 Research Objectives	6
1.5 Measure of Success	6
1.6 Evidence of Effectiveness	7
1.7 Primary Research Contributions	7
1.8 Thesis Structure	8
<i>Chapter 2. Literature Review</i>	<i>9</i>
2.1 Introduction	9
2.2 Artificial Intelligence (AI)	10
2.2.1 <i>What is AI?</i>	10
2.2.2 <i>Key AI Developments</i>	11
2.2.3 <i>State of the Art</i>	14
2.2.4 <i>Machine Learning</i>	15
2.2.5 <i>Deep Learning</i>	17
2.2.6 <i>The Future of AI and the Fourth Industrial Revolution</i>	18
2.3 Data Ethics	19
2.3.1 <i>Teaching Data Ethics</i>	22
2.4 Representation and Diversity	23
2.4.1 <i>Widening Participation</i>	24
2.5 Learning Theory	26
2.5.1 <i>The Concept of Learning</i>	26
2.5.2 <i>Theories of Learning</i>	28
2.5.3 <i>Learning Models and Strategies</i>	31
2.5.4 <i>Instructional Design Frameworks</i>	34
2.6 Computer Science Education	37
2.6.1 <i>Programming</i>	38
2.6.2 <i>AI, Machine Learning and Data Science</i>	40
2.7 Pedagogical Content Knowledge and Threshold Concepts	44
2.8 Cognitive Barriers	47
2.8.1 <i>Mathematics Anxiety</i>	47
2.8.2 <i>Confidence and Self-Efficacy</i>	51
2.8.3 <i>Metacognition</i>	53
2.8.4 <i>Self-Regulated Learners</i>	54
2.9 Online Learning	55
2.9.1 <i>Massive Open Online Courses (MOOCs)</i>	56

2.9.2 <i>Online Pedagogy</i>	57
Summary.....	58
Chapter 3. Methodology	61
3.1 Introduction.....	61
3.2 Data Sources	61
3.2.1 <i>Online Review of Modules</i>	62
3.2.2 <i>Questionnaires for Lecturers</i>	63
3.2.3 <i>Interviews with Lecturers</i>	65
3.3 Case Studies.....	66
3.3.1 <i>Pre-Module Questionnaires</i>	67
3.3.2 <i>Observation of Lectures</i>	70
3.3.3 <i>One Minute Paper Technique</i>	70
3.3.4 <i>Post-Module Questionnaires</i>	72
3.4 Analysis Methods.....	74
3.4.1 <i>Statistical Methods</i>	74
3.4.2 <i>Qualitative Methods</i>	78
3.5 Limitations of the Data	79
3.6 Summary.....	80
Chapter 4. Current AI Education Provision and Experiences Within Higher Education	81
4.1 Introduction.....	81
4.2 Study Context	81
4.3 Online Review of Modules.....	82
4.3.1 <i>Universities Under Review</i>	83
4.3.2 <i>Machine Learning Modules</i>	84
4.3.3 <i>Deep Learning Modules</i>	91
4.3.4 <i>Artificial Intelligence Modules</i>	91
4.3.5 <i>Data Science Modules</i>	93
4.3.6 <i>Summary of Findings</i>	94
4.4 Data on Lecturers' Experiences	96
4.4.1 <i>Questionnaires</i>	96
4.4.2 <i>Interviews</i>	97
4.4.3 <i>Key Themes</i>	100
4.5 Case Studies	101
4.5.1 <i>Newcastle University Machine Learning Module</i>	101
4.5.2 <i>Artificial Intelligence Module</i>	115
4.5.3 <i>Machine Learning and Computer Vision Module</i>	118
4.5.4 <i>Key Themes from Case Studies</i>	119
4.6 Discussion	121
4.6.1 <i>Current Provision</i>	121
4.6.2 <i>Lecturer and Student View</i>	123
4.6.3 <i>Educational Background</i>	124
4.6.4 <i>Learning and Teaching Strategies</i>	126
4.7 Summary.....	128
Chapter 5. An Online Learning Tool for an Introduction to Machine Learning.....	130
5.1 Introduction.....	130
5.2 Rationale for Creating an Online Tool	130
5.3 Tutorial Requirements	135

5.3.1 Tutorial Design Requirements	136
5.3.2 Tutorial Technical Requirements	137
5.4 Design	140
5.5 Numbas	141
5.5.1 How Numbas will be adapted	143
5.6 Implementation	148
5.7 Tutorials	150
5.7.1 Maths and Stats for Machine Learning	152
5.7.2 Data Pre-Processing	155
5.7.3 Overview of AI	157
5.7.4 Overview of Machine Learning	158
5.7.5 Machine Learning Algorithms	159
5.7.6 Overview of Deep Learning	161
5.8 Testing	163
5.8.1 R.1 Ability to host engaging multi-modality content	164
5.8.2 R.2 Capability for formative assessment	165
5.8.3 R.3 Provide learners with feedback	165
5.8.4 R.4 User interface must be attractive, consistent and responsive to the user	167
5.8.5 R.5 Tutorials must be interactive and engage users to maintain motivation	167
5.8.6 R.6 Ensure access to the tool is readily available	168
5.8.7 R.7 No new personally identifiable data should be collected through the tool	168
5.8.8 Expert Requirement Verification	169
5.9 Example of Use	169
5.10 Summary	170
<i>Chapter 6. Evaluation of MetaLearning</i>	172
6.1 Introduction	172
6.2 Comparison with Existing Tools	172
6.2.1 Evaluation of Numbas	173
6.3 Recruitment of Participants to Review MetaLearning	174
6.4 Student Questionnaire	174
6.4.1 Student Questionnaire Methodology	174
6.4.2 Student Questionnaire Analysis Methods	178
6.4.3 Student Questionnaire Results	179
6.4.4 Maths and Stats Knowledge Survey Results	190
6.4.5 Student Results Summary	193
6.5 Lecturer Interviews	194
6.5.1 Lecturer Interview Methodology	194
6.5.2 Lecturer Interview Analysis Methods	196
6.5.3 Lecturer Interview Results	197
6.6 UKICER 2021 Workshop	200
6.6.1 Running the Workshop	200
6.6.2 Findings	203
6.7 Discussion	206
6.8 Summary	209
<i>Chapter 7. Towards a Framework for Teaching Machine Learning</i>	210
7.1 Introduction	210
7.2 Barriers to Learning Machine Learning	211
7.2.1 Educational Background and Mathematics Anxiety	212

7.2.2 Metacognition, Self-Efficacy and Self-Regulation	216
7.2.3 Threshold Concepts and Pedagogical Content Knowledge	219
7.3 Online Resource for Introduction to Machine Learning (MetaLearning)	221
7.3.1 Online Pedagogy	222
7.3.2 Framework of Topics for an Introductory Machine Learning Course	223
7.3.3 MetaLearning Review	226
7.4 Guidance to Aim to Improve Student Satisfaction and Attainment	226
7.4.1 Module Pre-Requisites and Pedagogy	227
7.4.2 Widening Participation	230
7.5 Summary	231
<i>Chapter 8. Conclusion</i>	232
8.1 Review of Research Aim and Questions	232
8.2 Review of Research Objectives	233
8.3 Research Contributions	234
8.4 Limitations of the Work	236
8.5 Future Work	237
8.6 Conclusion	239
<i>References</i>	242
<i>Appendix A: Online Questionnaire for Lecturers</i>	278
<i>Appendix B: Interview Questions for Lecturers</i>	279
<i>Appendix C: Pre-Module Student Questionnaire</i>	280
<i>Appendix D: Observation Guide</i>	283
<i>Appendix E: Week 2 One Minute Paper</i>	287
<i>Appendix F: Student Post-Module Questionnaire</i>	288
<i>Appendix G: Miscellaneous Machine Learning Modules</i>	289
<i>Appendix H: Assessment from AI Module</i>	291
<i>Appendix I: Links to MetaLearning Tutorials</i>	294
<i>Appendix J: Student Questionnaire – MetaLearning Review</i>	295
<i>Appendix K: Maths and Stats Knowledge Survey</i>	296
<i>Appendix L: Information Sheet for Lecturers Reviewing MetaLearning</i>	297
<i>Appendix M: History of AI Timeline from MetaLearning</i>	299

Table of Figures

Figure 1: Overview of AI, Machine Learning and Deep Learning (based on (Jeffcock, 2018; Ceron, 2019)).....	1
Figure 2: The percentage of underrepresented students enrolled on the postgraduate conversion course (Office for Students, 2021a).	24
Figure 3: The main aspects of learning (based on (Illeris, 2018))	27
Figure 4: Kolb's Experiential Learning Cycle	31
Figure 5: Changes to the Bloom taxonomy from Krathwohl and Anderson (image altered from (Moseley <i>et al.</i> , 2005)	36
Figure 6: An example of one of the tutorials from (Code.org, 2021a)	42
Figure 7: Questions 3d and 3e from pre-module questionnaire	69
Figure 8: The One-Minute Paper (Week 1).....	72
Figure 9: The educational level at which Machine Learning modules are offered.....	85
Figure 10: Previous degree of students from Newcastle University	103
Figure 11: Most frequent responses to Q1: Which topic(s) did you find the most difficult this week?	112
Figure 12: Q2: Which topics are you still unsure on?	113
Figure 13: E-Learning Usability Framework (adapted from Zaharias and Poylymenakou, 2009)	138
Figure 14: Introductory page for the tutorials.....	145
Figure 15: Design of the individual tutorial pages	146
Figure 16: Message upon selection of 'End Exam'	147
Figure 17: Performance summary	147
Figure 18: Overview of AI tutorial screen	151
Figure 19: Introductory screen for the Overview of AI tutorial	151
Figure 20: Summary of marks from the Maths and Stats tutorial	155
Figure 21: Example of a multiple-choice question relating to variables.....	156
Figure 22: Low stakes question pertaining to mapping of a dataset.....	156
Figure 23: Visualisation of the disciplines which contribute to the field of AI	157
Figure 24: Visualisation of the differing subdomains within AI and how these intersect	158
Figure 25: A visual reference guide to the different types of learning algorithms.....	159
Figure 26: Visual representation of a Decision Tree.....	160
Figure 27: Visualisation of the steps for implementing the K-Means algorithm	161
Figure 28: Graphical representation of gradient descent.....	162
Figure 29: Example of feedback of a correct answer	165
Figure 30: Example of feedback with an advice section	166
Figure 31: Options to change the visual display	167
Figure 32: Question 7 of MetaLearning review questionnaire	176
Figure 33: Usability questions from the MetaLearning questionnaire	177
Figure 34: Question on features which can be improved on	177
Figure 35: Question relating to confidence in applying knowledge in a practical context....	178
Figure 36: Student gender demographics from the student MetaLearning questionnaire ..	181
Figure 37: Student views on whether MetaLearning improved confidence	182
Figure 38: Responses relating to confidence in mathematics	183
Figure 39: Frequency response to knowledge survey question	184
Figure 40: Percentage response to the Q: Which of the tutorials did you find the most difficult?.....	185

Figure 41: Responses to the questions relating to the low-stakes quizzing strategy.....	186
Figure 42: Frequency response to the question ‘On a scale of 1-5 how easy was the system to use?’	187
Figure 43: Frequency response indicating how confident respondents are applying their knowledge in a practical context	189
Figure 44: Interview questions about MetaLearning	196
Figure 45: Completed padlet for exercise one	203
Figure 46: Results from the padlet vote from exercise three	205
Figure 47: Diagrammatic Overview of Framework	211

Table of Tables

Table 1: Skills that graduates are lacking (from GOV.UK (2021))	40
Table 2: The different characteristics of students with high and low self-efficacy (Source: Hattie, 2012).....	52
Table 3: Proposed analysis techniques.....	74
Table 4: Overview of the Introduction to Machine Learning modules.....	87
Table 5: Overview of the Machine Learning modules	88
Table 6: Overview of Advanced Machine Learning modules.....	90
Table 7: Overview of undergraduate Artificial Intelligence modules	93
Table 8: List of all topics mentioned by participants within the interview.....	98
Table 9: Age demographics for students on Machine Learning module	102
Table 10: Gender for students on the Machine Learning module.....	103
Table 11: Confidence in Maths ANOVA	104
Table 12: Gender and programming skill level crosstabulation	105
Table 13: Responses to the description of Machine Learning	106
Table 14: Learning strategies students propose to use in module	108
Table 15: Topics from the observed sessions of the Artificial Intelligence module	116
Table 16: Commonly taught topics identified in Chapter 4	133
Table 17: Framework for critical elements of online learning settings (altered from (Oliver, 2001))	149
Table 18: The question groups and overview of content of the Maths and Stats tutorial....	154
Table 19: Introductory AI Courses Comparison Table	173
Table 20: Participants level of study UG (Undergraduate), PG (Postgraduate).....	180
Table 21: Responses relating to improvement in Maths and Stats	183
Table 22: Results from the question: how likely are you to reuse the tool?	188
Table 23: All questions from the knowledge survey.....	192
Table 24: The questions which received the lowest overall confidence rating	193
Table 25: Workshop exercises and questions	202
Table 26: Topics mentioned in lecturer questionnaires and interviews.....	225
Table 27: Topics identified as most commonly taught within this domain	225

Chapter 1. Introduction

1.1 Research Context

Artificial Intelligence (AI) endeavours to comprehend the concept of intelligence, what this constitutes, as well as to create and build intelligent systems (Russell and Norvig, 2013). There are a number of domains within the field of AI (as shown in Figure 1) as well as specific application areas such as robotics and computer vision. Machine Learning, a sub-domain within AI is “a branch of artificial intelligence that allows computer systems to learn directly from examples, data, and experience” (The Royal Society, 2017, p.5). It is this capability which enables computers to perform tasks by learning from data instead of using pre-programmed rules. A sub-domain of the field of Machine Learning, is Deep Learning which “allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction” (Lecun, Bengio and Hinton, 2015, p.436). Due to the phenomenal successes of Deep Learning within application areas such as computer vision, Deep Learning has contributed to the recent resurgence of AI which has increased the uptake of this technology within industry and the popularity of this subject amongst students. Within this thesis, for the sake of clarity, AI will be used to denote courses encompassing Machine Learning and Deep Learning, without loss of generality.

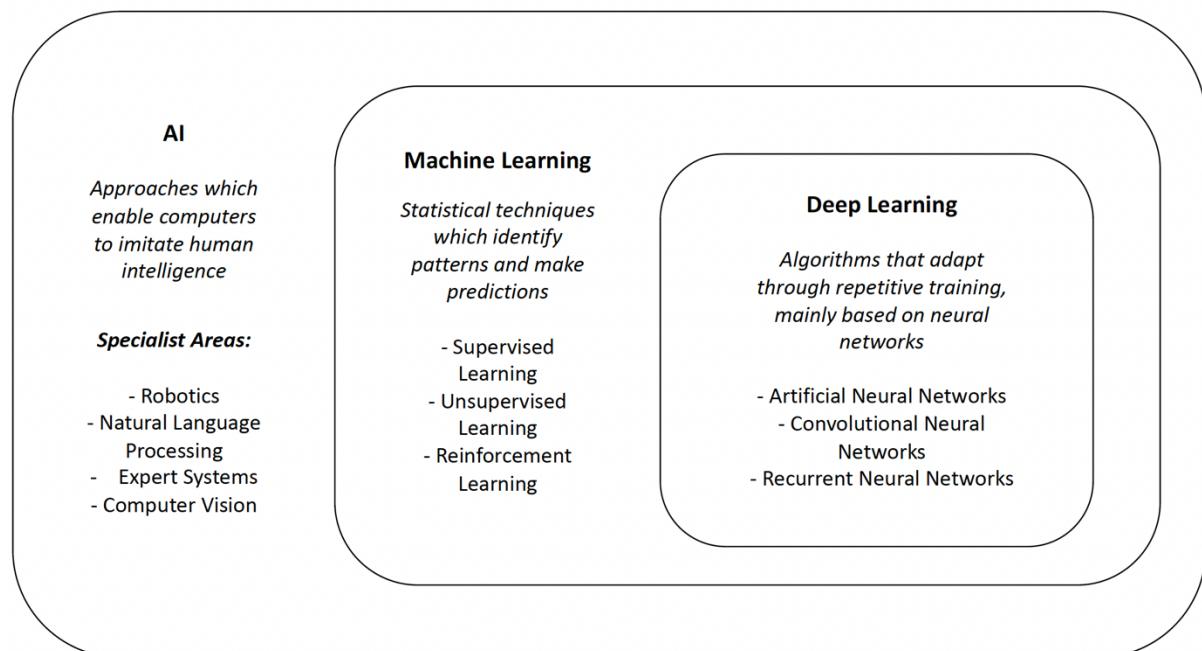


Figure 1: Overview of AI, Machine Learning and Deep Learning (based on (Jeffcock, 2018; Ceron, 2019))

Individuals skilled within the AI domain are highly sought after within industry due to the current skills shortage which is becoming an increasing problem for recruiters as the UK digital economy grows (Department for Digital Media and Sport, 2017). This trend is also experienced worldwide with a global shortage of workers with skills and experience within this domain (IBM, 2020a). AI is a complex area and self-study for existing employees is unlikely to fill the skills gap. Therefore, we need more students graduating with AI skills. Many educational institutions are now offering some form of an AI course due to the rising interest in study at higher education (HE) level. In 2017, within the UK there were twenty-six universities offering undergraduate courses in AI, with more than thirty graduate programmes running across twenty universities (HM Government, 2018). However, more recent findings in 2021 suggest that there is currently a demand for an expansion to places on such courses of a “five to tenfold increase” (UK AI Council, 2021). The number of students taking courses within this domain would require “significant increases in numbers at both levels [undergraduate and postgraduate]” (Hall and Pesenti, 2017, p.54) to fill the UK deficit in individuals skilled within this domain. A similar trajectory is seen in the US and European Union (Stanford University, 2021). 93% of organisations within the US and the UK advise that they consider AI to be a business priority, however 51% acknowledge that they do not have appropriate staffing levels to implement their AI strategies (Bourne, 2019).

Two issues which may impact on attracting and retaining students within this topic are:

1. The best practices in pedagogy to teach the specialist skills required are still relatively unexplored.
2. AI is an advanced subject that combines aspects from both Mathematics and Computer Science where students may lack confidence in their ability to do well in such courses, especially if they suffer from mathematics anxiety or have a general lack of confidence in their technical skills.

Mathematics anxiety is associated with “feelings of tension and anxiety that interfere with the manipulation of numbers and the solving of mathematical problems” (Richardson and Suinn, 1972). The level of mathematics and Computer Science knowledge required to understand and execute AI is often not taught within a single course. The level of both computing and mathematical skills required by the student is determined by the type of AI module they are undertaking. There are two potential strands of AI education:

- How to innovate and create new AI methodologies.
- How to understand and apply existing AI/Machine Learning algorithms and techniques.

The prerequisites for an AI module may differ based upon the fundamental aspects of the course, with the creation and improvement of new AI methodologies requiring greater in-depth mathematical knowledge. Therefore, there is potential disparity within AI cohorts in relation to educational background.

1.2 Research Motivation

As outlined in the previous section there is a growing demand for individuals with skills and experience within the AI domain, alongside an increase in interest and educational provision at higher educational level. However, there is a current lack of existing research pertaining to recommended education practices and appropriate pedagogy for teaching this domain which requires a high level of both mathematical and computational knowledge.

Within this thesis qualitative research and existing studies are used to determine the opinion of AI within the scholarly community and to determine any misconceptions prospective AI students may have relating to this domain. Identifying any misinterpretations of this topic, particularly with learners, will help to identify any mental models individuals may already have. A mental model can be considered a “mental simulation” of the situation or problem with the aim of assisting its builder to “explain and make predictions” about the represented situation (Greca and Moreira, 2000, p.3). Gaining an in-depth understanding of opinions/impressions of AI could potentially lead to the development of strategies to encourage a larger number of people to become skilled in this domain – thus accommodating the rising demand. This thesis aims to provide findings relating to the two issues identified in the previous section in relation to attracting and retaining students and determine the current experiences of learners studying this domain and how this can be improved.

Surveys completed by Cameron and Maguire (2017) and Ipsos MORI (2017) indicate that individuals who could be considered “digital natives”, a definition formulated by Prensky (2001) to signify individuals who grew up with computer and internet access, are more familiar with the term Machine Learning. However, the studies also found that the participants were

not especially interested in how Machine Learning worked, “in part due to the complexity of the technology being something they assumed they would not be able to understand” (Cameron and Maguire, 2017). The assumption that Machine Learning is too complex to comprehend, especially from digital natives, may be an issue relating to the recruitment of students to modules/courses within this domain. It may also be an issue for mature students or employees who wish to reskill in this domain who are less familiar with the term and are not deemed digital natives.

Determining the educational background, particularly the prior mathematics attainment level and programming experience of students on AI courses could potentially determine whether there are any topic or skill gaps which may lead to difficulties when completing these courses. Alongside investigation of AI, Data Science courses are also reviewed. This field is “still in its formative period” (Young, Wajcman and Sprejer, 2021) with a definitive outline of the tools and methods within this discipline to be determined. However, many data scientists are proficient in Machine Learning (Fayyad and Hamutcu, 2020). Therefore, any existing research pertaining to Data Science education may be pertinent to this study.

Identifying some of the potential barriers students encounter when undertaking education within the AI domain is a central tenant of this research. Of particular focus are issues of confidence and self-efficacy and how these might impact a student’s willingness or ability to learn. Steven and Thomas (2019, p.28) advise that “there will rarely be only one barrier facing a particular group,” instead the students may potentially be encountering a number of barriers which may amalgamate. Prior research pertaining to barriers to learning often use a deficit model where the underrepresented group are charged with overcoming these issues, “rather than assessing the impact of institutional infrastructure, entry requirements, course structure and student experience” (Steven and Thomas, 2019, p.5). Relating to the concept of self-efficacy is metacognition which is the “ability to articulate and regulate the mental processes that we use to construct our knowledge understanding and skills” (Luckin, 2018, p.43). Metacognition, self-efficacy and self-regulation are all concepts which have been identified to “help students to organise their study activity independently and effectively” (Cera, Mancini and Antonietti, 2014). These specific cognitive skills are a key focus in this research as they can assist learners in persisting through difficulties such as threshold concepts, which are core concepts of the domain which, once understood can lead to a

different view of the subject area (Kiley and Wisker, 2009), and other barriers identified and investigated within this thesis.

Identification of the threshold concepts and pedagogical content knowledge within the AI domain has yet to be established. Therefore, this research will initialise this process as it is believed that identifying the threshold concepts will enable greater teacher understanding of the specific topic which can cause students difficulty and will help to guide future learning design and best practice. Although threshold concepts can cause students difficulty, they can also lead to greater understanding of key ideas within the field of AI if taught effectively. Walker (2013) advises that threshold concepts can be regarded as a “particular state of expert knowledge” and that they are often the parts of a module where students ‘get stuck.’ For students who have not yet fully comprehended a threshold concept, they may learn new topics in a more disjointed fashion as they cannot yet integrate this new concept into their current way of thinking. Students who have more sophisticated metacognition skills will be better equipped to navigate through the threshold concepts as metacognitive processes are “associated with enhanced cognitive performance” (Luckin, 2018, p.45). Once a student has comprehended the threshold concept, they can then integrate different aspects of the overall subject into their analysis of problems (Land *et al.*, 2005). Preparing for threshold concepts within courses should ensure that lecturers can implement strategies to assist students when they start learning these concepts, thus demonstrating that they can tolerate learner confusion (Cousin, 2006) and also helping the lecturers formulate differing approaches to better student understanding. Students often have a non-linear, complex path towards learning which, unless communicated or comprehended by the lecturer can often lead to miscommunication relating to student progress (Lucas and Mladenovic, 2007).

1.3 Research Aim and Research Questions

The aim of this research is to determine the barriers students face when learning AI as well as the difficulties encountered by lecturers teaching this domain to initiate a framework of best practice for AI education.

The research aims to address the following research questions (RQ):

RQ1: What is considered good practice relating to the teaching of AI?

RQ2: What are the current perceived difficulties experienced by both students and lecturers relating to AI?

RQ3: How do cognitive mitigation strategies alleviate any identified issues encountered by students learning this domain?

The uniqueness of the research questions stem from the current lack of pedagogical research relating to AI education. Identification of the barriers to learning will be considered through a combination of differing methodologies to determine both student and lecturer perspective. Focusing on cognitive mitigation strategies to assist learners in overcoming the encountered difficulties also offers a singularity to this research, as there are currently limited studies relating to use of these methods to aid learning in Computer Science education.

1.4 Research Objectives

The research objectives (RO) of this work are as follows:

1. To identify and examine the barriers that might impact upon student attainment in Machine Learning courses, using Machine Learning modules at Newcastle University and other participating institutions.
2. To use the results to:
 - a. Identify the threshold concepts within Machine Learning.
 - b. Create a learning resource tool which aims to assist students within their learning of Machine Learning through tutorials and visualisation of identified threshold concepts and Machine Learning topics.
 - c. Improve student's metacognition and self-regulation regarding their learning of Machine Learning through implementation of strategies such as testing as a learning tool and knowledge surveys within the learning resource.
 - d. Create a framework of topics for an introductory Machine Learning course.
 - e. Discover ways of improving student satisfaction and attainment within these courses.

1.5 Measure of Success

The criteria for success pertaining to the aims, questions and objectives of this research include clear responses from the varying data collection methods. Relating to RO2.b and RO2.c

receiving positive feedback from both students and practitioners relating to the learning resource will be a measure of success alongside continuous use of the resource upon study completion. Defining some best practices for education within this domain will also contribute to the initialisation of the framework, culminating in some of the findings being published.

1.6 Evidence of Effectiveness

Evidencing the effectiveness of the approaches used to achieve the research aim and objectives will be evaluated against existing studies which employ a similar methodology. The combination of data collection methods within this study have not been employed previously for education research purposes, therefore findings relating to the specific methods will be evaluated. The findings will also be compared to other outcomes within the Computer Science education domain.

1.7 Primary Research Contributions

There are a number of contributions from this thesis, the main contribution is identification of the barriers and issues students face when completing a module within the AI domain. These findings stem from the online review of modules to determine the current education provision within this domain as well as statistical analysis of data collected through qualitative methods such as questionnaires, interviews and observation carried out on modules within the AI domain. Although these findings were derived specifically from AI modules, the findings may be pertinent to other domains within Computer Science.

As well as identification of difficulties and obstacles related to learning the topic of AI, several mitigation strategies have been trialled within this research to determine their effectiveness in assisting students in overcoming the identified barriers. These findings can also be generalised to other educational scenarios where learners may lack confidence or self-efficacy in their skills, and where methods to improve their metacognition can assist them to become more self-regulated learners.

The learning resource created as part of this study encapsulates a range of the best practices identified through the findings from the qualitative research. The framework of topics created in fulfilment of RO2.e. forms the basis for the content for the resource alongside the

preliminary establishment of the threshold concepts. The learning resource offers a comprehensive introduction to the AI domain, with specialisation in the field of Machine Learning.

1.8 Thesis Structure

Chapter 2 of this thesis provides a review of the pertinent background literature used within this research (RQ1, RO1). Within Chapter 3, the methodology for the analysis of the current education provision in AI is discussed, including the data sources and differing data collection methods as well as the statistical methods used to evaluate the data from this strand of the study. Chapter 4 details the results from the analysis carried out and summarises the findings and the implication for AI education provision (RQ1, RQ2, RO1, RO2.a). Chapter 5 details the creation of the learning resource including the design, content and inclusion of the mitigation strategies outlined in the research objectives (RO2.b, RO2.c, RO2.d). Chapter 6 outlines the methodology for the review of the learning resource as well as the analysis of the findings, including the lecturer and student view (RQ1, RQ2, RO1, RO2.a, RO2.c). Chapter 7 presents all of the results from this research and presents these as an initialisation of a framework of best practice for teaching this domain (RQ1, RO2.e). Chapter 8, the conclusion, includes a summary of the work completed as well as an insight into further work which will be carried out. A full set of references and appendices are provided.

Chapter 2. Literature Review

2.1 Introduction

As discussed in Chapter 1, Machine Learning is now considered an integral aspect of Computer Science (Shapiro, Fiebrink and Norvig, 2018) and the popularity of courses within this domain are rapidly increasing alongside the need for individuals specialised within this area. This is due to the current skills shortage which is predicted to rise apace with the growth of the digital economy (Department for Digital Media and Sport, 2017), both within the UK and globally. There is also a need for more Artificial Intelligence (AI) based courses, including MOOCs (Massive Open Online Courses) and continuing professional development courses to decrease the skill shortage and to increase the number of people trained with these specialist skills (Hall and Pesenti, 2017). However, there is currently a lack of research specifically relating to the best educational practices within this domain. With the increasing demand for graduates who are skilled particularly in Machine Learning, programming and data ethics (Department for Digital Media and Sport, 2021), and increasing numbers of students partaking in these courses, it is important to identify the best pedagogic practices for teaching within this discipline and the barriers students may face when participating in such courses.

This chapter reviews the literature pertaining to AI and higher education, specifically different learning theories, threshold concepts and the best practices for Computer Science education. Higher education can be defined as education provided within a postsecondary institution such as a university which usually leads to a named degree, diploma or certificate of higher studies upon course completion (Editors of Encyclopedia Britannica, 2013). This literature review also examines cognitive barriers which students may encounter such as mathematics anxiety as well as issues pertaining to cohort diversity and AI. Literature relating to online learning has also been critiqued to determine the current conventions relating to this form of instruction and how to relate this to the barriers students face as a solution to alleviate the current inequity of provision.

Within this research there is an overlap between some aspects of domain terminology, for example, Deep Learning within the context of AI means learning “in a form of multiple levels of representation and abstraction to make up higher level information from lower level information” (Zhang *et al.*, 2018). Whereas deep learning in an educational context entails “a

deep approach to learning” which “involves an intention to understand and impose meaning” (Smith and Colby, 2007). Therefore, inclusion of this terminology will include the context in which this term is situated to avoid confusion.

2.2 Artificial Intelligence (AI)

Within this section, important aspects of the Artificial Intelligence domain will be covered, including an overview of what AI is, a brief history of this domain to provide context to the current state of the art alongside a discussion of particular subdomains pertinent to this study, including Machine Learning and Deep Learning. The future of AI will also be discussed to identify how the outcomes of this research may have potential impact and relevance to this fast-paced domain.

2.2.1 What is AI?

Artificial Intelligence (AI) “attempts not just to understand but also to build intelligent entities” (Russell and Norvig, 2013, p.1). The term *Artificial Intelligence* was conceived of in 1956 at the official birthplace of the field, Dartmouth College, by John McCarthy. Alongside colleagues such as Marvin Minsky, McCarthy hosted a two-month project relating to the study of AI. The project was based on “....the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.” (McCarthy *et al.*, 2006). The aim is to construct machines which display behaviour which, if observed in humans, would be described as “intelligent” (Feigenbaum, 1963).

One of the prominent aims within the field of AI is to build a machine which has common sense and programs which learn from their experiences as efficaciously as humans (McCarthy, 1959; Minsky, Singh and Sloman, 2004). Researchers worked with the hypothesis that “human thinking is wholly information-processing activity” and that these processes are explainable and understandable (Feigenbaum, 1963).

AI evolved from its early beginnings into a field which is integrated into a range of differing sub-fields due to its universal relevance in any intellectual tasks (Russell and Norvig, 2013). In the subsequent years after the inception of the field of AI, focus of this domain centred on specific techniques which can handle a certain class of tasks, which however do not generalise

well (Minsky, Singh and Sloman, 2004). These included the General Problem Solver, Machine Translation and Microworlds (Russell and Norvig, 2013) which will be discussed in greater detail in the following sections. The subfields within AI of Machine Learning and Deep Learning gained prominence as a result of the breakthroughs within these specific tasks. Overall the AI domain has vastly changed since its inception as a consequence of the numerous interpretations and perceptions of what intelligence is and what it entails (Martínez-Plumed *et al.*, 2018).

2.2.2 Key AI Developments

As previously discussed, Artificial Intelligence (AI) was defined by John McCarthy in 1956 alongside colleagues Marvin Minsky, Nathaniel Rochester and Claude Shannon as they hosted a two-month study. Although the study in 1956 was the first usage of the term AI, work done by Warren McCulloch and Walter Pitts in 1943 is now recognised as the first work in AI. McCulloch and Pitts drew inspiration from the function of neurons within the human brain, propositional logic and Turing's theory of computation to propose a model of artificial neurons and networks which could learn (Russell and Norvig, 2013). Artificial neural networks are a key focus within the subdomain of Deep Learning (as discussed in Section 2.2.5).

One of the most influential suppositions of AI came in 1950 when Alan Turing published his article “Computing Machinery and Intelligence”. In this seminal paper, Turing proposed the idea of the child programme in which a programme is proposed which simulates a child's mind which then can be “subjected to an appropriate course of education” to simulate an adult's mind (Turing, 1950). The influence of the understanding and idea of human cognition is intrinsically linked with the field of AI and has been incorporated into this field since inception. Within the paper by Turing, he also detailed the imitation game, which is often referred to as the Turing Test, although there is some controversy that these are indeed two separate tests (Sterrett, 2000). The Turing Test attempts to assess the “machine's ability to imitate a human being” (Saygin, Cicekli and Akman, 2000), this encompasses a human interrogator attempting to distinguish between a computer and a human subject based on their various replies to questions posed by the interrogator (Encyclopedia Britannica, 2021).

Within Minsky's work at MIT within the 1960s, he supervised a number of students "who chose limited problems that appeared to require intelligence to solve" (Russell and Norvig, 2013, p.19), similar to the current AI trajectory. These domains were known as microworlds. Microworlds became very popular within the field of education as they allowed students to explore a particular domain whilst conforming to the laws and constraints of the subject matter (Miller, Craig, Lehman, Jill and Koedinger, Kenneth, 1999). One of the most influential microworlds was the blocks world which was used for the understanding natural language program from Winograd in 1972.

The 1980s signified a change in that AI became a commercial viability instead of a purely academic pursuit, the R1 became the first successful commercial expert system which saved the company, Digital Equipment Corporation an estimated \$40 million a year (Russell and Norvig, 2013). Many companies followed suit and invested money in research and development related to AI, however many companies failed to deliver on their aspirations which led to many companies abandoning such projects and contributing to the "AI Winter" (Russell and Norvig, 2013).

In 1988, Judea Pearl published *Probabilistic Reasoning in Intelligent Systems*, in which he suggests the use of probability as an "initial model of human reasoning" (Pearl, 1988). Pearl explained that probabilistic formalisms enable us to "summarize the presumed existence of exceptional conditions without explicating the details of their interactions unless the need arises." The Bayesian formalism was also constructed during this period which provides "a formalism for reasoning about partial beliefs under conditions of uncertainty" (Pearl, 1988). This highly influential work by Judea Pearl led to a "new acceptance of probability and decision theory in AI" (Russell and Norvig, 2013, p.26).

The backpropagation algorithm was first conceived of in the 1960s, however it was reinvented by a number of different groups in the 1980s (Russell and Norvig, 2013). One of the most influential uses of the backpropagation algorithm was within neural networks as introduced by Rumelhart, Hinton and Williams in their 1986 paper *Learning representations by back-propagating errors* (Rumelhart, Hinton and Williams, 1986). Backpropagation is a central component of most artificial neural networks, therefore it is key for developers to have a level of understanding of the fundamentals of this process.

The 1990s saw a number of algorithms come into fruition, the Random Forest algorithm was created in 1995 as a method to expand upon Decision Trees to increase accuracy for training and unseen data, this is achieved by building multiple trees “in randomly selected subspaces of the feature space” (Ho, 1995). In the same year Cortes and Vapnik published a paper detailing Support Vector Machines and their use for “two-group classification problems” (Cortes and Vapnik, 1995). Both of these machine learning algorithms are widely used and applied today.

One of the most high-profile events in AI occurred in 1997 when the IBM Deep Blue computer beat the world chess champion, Gary Kasparov. The Deep Blue computer achieved this victory through custom circuits, parallel search engines and various search algorithms (Hsu, Campbell, Murray and Hoane, 1995). The architecture used within the Deep Blue went on to be used for various other tasks such as financial modelling, data mining and molecular dynamics (IBM, 2020b). Deep Blue is considered an important milestone in AI as it “provoked considerable thought on the subject of what intelligence is all about” (Newborn, 2000).

Some of the most challenging goals within the field of Artificial Intelligence are centred around algorithms which learn without any previous knowledge to gain proficiency in complicated domains, such as the game of Go (Silver *et al.*, 2017). In 2015, AlphaGo became the first computer system to beat a world champion at Go, the second iteration of this program defeated Lee Sedol in 2016, who has won 18 international titles in Go (Silver *et al.*, 2017). The most recent iteration of AlphaGo, entitled AlphaGo Zero utilises a form of reinforcement learning, AlphaGo Zero is more powerful than the previous versions in that “it is able to learn *tabula rasa*” (Silver and Hassabis, 2017). Achieving this capability has been a persisting goal within the field of Artificial Intelligence.

As outlined in this section, AI was established as an academic discipline in 1956. The key milestones of this discipline including the earliest examples of neural networks helped shape the Deep Learning subdomain, including concepts such as backpropagation which are core aspects of the discipline today. The commercialisation of AI within the 1980s highlighted industry demand for this technology, which is increasing alongside the continuous breakthroughs within this domain.

2.2.3 State of the Art

The field of Artificial Intelligence is a fast-moving domain, where new innovations are quickly being discovered and deployed. The availability of vast datasets and the advances in computing power and GPUs from companies such as NVIDIA (2020) have had particular influence within the subfield of Deep Learning and has led to an AI resurgence. A number of prevalent libraries for creating Machine Learning models, such as TensorFlow (Abadi *et al.*, 2015) and Keras (Chollet, 2015), have also widened the scope of this field in that people now have the tools and resources to implement their own AI projects and develop their skills and knowledge within this field.

The AI resurgence has been spearheaded by the advances within Deep Learning. As recognition of the effect that this subgenre has had on the field of AI the 2018 Turing Award was presented to three of the most eminent members of this field. Yoshua Bengio, Geoffrey Hinton and Yann LeCun received this award in recognition of their “conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing” (ACM, 2018). Individually, they have contributed key breakthroughs within Deep Learning, Hinton developed the idea of Boltzmann machines (Fahlman, Hinton and Sejnowski, 1983), Bengio and his work on Generative Adversarial Networks (GANs) (Goodfellow *et al.*, 2020) and LeCun’s seminal work on Convolutional Neural Networks (CNNs) (LeCun *et al.*, 1998). All three work within the intersection of Neuroscience/Cognitive Science and Machine Learning (ACM, 2018).

Although breakthroughs are consistently being made in specific domains on individual tasks, there is a call for more focus on Artificial General Intelligence (AGI) and what Stuart Russell explains as “beneficial machines”, in which the machine is expected to achieve objectives outlined by the human based on our preferences (Russell, 2019). There is also an urgent need for greater transparency and explainability within the field of AI which has been demonstrated by a number of adversarial systems and the risks associated with these such as discrimination, opaque decision making and use for criminal purposes (European Commission, 2020). These concepts will be discussed in greater detail in Section 2.3 Data Ethics.

2.2.4 Machine Learning

There are a number of sub-disciplines which build upon the foundations of AI and one of these is Machine Learning. Machine Learning is “a branch of artificial intelligence that allows computer systems to learn directly from examples, data and experience” (The Royal Society, 2017). This capability enables computers to perform tasks by learning from data instead of using pre-programmed rules.

Machine Learning algorithms have been applied in a wide variety of domains including virtual personal assistants such as Siri and Alexa and in areas such as anomaly detection for example in identifying credit card fraud and machine translation (Goodfellow, Bengio and Courville, 2016). Some of the most common Machine Learning algorithms include Linear Regression (Stanton, 2001), Logistic Regression (Cramer, 2002), Decision Trees (Quinlan, 1986) and Support Vector Machines (Noble, 2006).

The core principle of Machine Learning is designing algorithms that “automatically extract valuable information from data” (Deisenroth, Faisal and Ong, 2020). Therefore, having a baseline understanding of the underpinning mathematical and statistical aspects of this domain is an essential aspect of becoming a practitioner within this area. For example, data within Machine Learning algorithms is represented as vectors and matrices which requires an understanding of linear algebra and matrix decomposition. Analytic geometry is also a fundamental element of Machine Learning to “formalize the idea of similarity between vectors” (Deisenroth, Faisal and Ong, 2020). Key to two of the central tenets of Machine Learning, dimensionality reduction and density estimation, is the foundational knowledge of probability and distributions. Probability theory enables us to introduce predictors and quantify levels of uncertainty. Cognition of these fundamental principles upon which Machine Learning algorithms are built can enable insight into the limitations of the Machine Learning algorithms as well as facilitate the implementation of new solutions within Machine Learning.

Chollet (2018) identified a universal workflow of Machine Learning, which is also pertinent to the sub-domain of Deep Learning, the workflow is based around 7 steps:

1. *Defining the problem and assembling a dataset:* this stage includes determining what the input data will be and identifying what type of problem you are trying to solve. For example is it a binary classification problem or is clustering required? Determining the problem will guide the choice of model, architecture and loss function.
2. *Choosing a measure of success:* it is important to identify what is meant by success for your model, for example is it precision, accuracy, recall or something more specific such as customer retention rate? The metric for success will also help guide the choice of loss function which is a method for “evaluating how well a specific algorithm models the given data” (Parmar, 2018).
3. *Deciding on an evaluation protocol:* once the success measure is defined, identifying an evaluation protocol is the next stage to determine how you’ll measure progress. Common evaluation protocols include K-fold cross validation and a hold-out validation set (Bengio and Grandvalet, 2004).
4. *Preparing your data:* this step involves formatting the dataset to ensure it is in an appropriate state to be input into the Machine Learning algorithm. Processes involved in this step include handling missing data and feature engineering.
5. *Developing a model that does better than a baseline:* at this stage the model is beginning to be trained. It is optimal to start with a small model which can beat a simple baseline, this is important to achieve statistical power. Once the baseline has been achieved, the activation function and loss function need to be determined. The activation function is pertinent to ANNs, it takes as input the output signal from the previous cell and “converts it into some form that can be taken as input to the next cell” (Jain, 2019). The loss function was discussed in step 2.
6. *Scaling up your model:* once statistical power has been achieved it is important to consider whether the model is sufficiently powerful. This point is of importance in Deep Learning, for example does the model have enough layers and parameters to properly model the problem.
7. *Regularizing your model and tuning your hyperparameters:* this step involves iterative modification to the model, including training and evaluating it on the validation data until it is as effective as possible. Potential methods to try include adding dropout, which is a regularization approach in neural networks which helps to prevent the issue of overfitting by “ignoring” units (i.e. neurons) during the training phase of a particular set of neurons which are chosen at random (Budhiraja, 2016). Iteration of feature

engineering can also be trialled to identify any features which don't appear to be as informative as well as trying different hyperparameters such as the number of units per layer to find the most favourable configuration.

(Chollet, 2018)

There are a number of issues to consider when creating a Machine Learning model, including underfitting and overfitting. Overfitting occurs when the model cannot generalize or fit well on unseen data, this can happen when the function corresponds too closely to the dataset (Tripathi, 2020). Underfitting is the opposite, this occurs when the model cannot model the training dataset or generalize to a new dataset, this is often easier to detect than overfitting (Tripathi, 2020). Another issue Machine Learning developers need to be cognisant of is the bias/variance tradeoff. A model with high-bias is most likely to underfit the training data, whereas a model with high-variance will overfit the training data. Géron (2017) advises that "increasing a model's complexity will typically increase its variance and reduce its bias." However, conversely, reducing the model's complexity will increase the bias and reduce its variance. Hence why the problem is called the bias/variance tradeoff.

2.2.5 Deep Learning

Deep Learning has gained precedence within AI over the past few years due to its prominence within research and practical applications. Yoshua Bengio, one of the pioneers within this field describes Deep Learning as an extension of earlier work on neural networks and as another approach to Machine Learning but one which is influenced by the brain (Bengio and Ford, 2018). The main advantage of Deep Learning is that it is "very good at discovering structures in high-dimensional data and is therefore applicable to many domains of science, business and government" (Lecun, Bengio and Hinton, 2015).

Prominent Deep Learning models include Convolutional Neural Networks (CNN) (LeCun *et al.*, 1998), Recurrent Neural Networks (RNN) (Hochreiter and Urgen Schmidhuber, 1997) and Generative Adversarial Networks (GAN) (Goodfellow *et al.*, 2020). These algorithms have been applied in tasks such as image recognition, speech recognition and prediction tasks in a variety of domains such as identifying activity of drug molecules and gene expression (Lecun, Bengio and Hinton, 2015).

Comparable to Machine Learning algorithms, of which Deep Learning algorithms are a sub-domain, is the importance of an understanding of key mathematical concepts to understand the inner workings of these models. The gradient descent algorithm is pertinent to have an understanding of to comprehend how to minimise the loss/cost function. Gradient descent is prominent within both Deep Learning and Machine Learning algorithms functioning as a “iterative first-order optimisation algorithm” (Kwiatkowski, 2021) which is employed to determine the local minimum or maximum of a given function. However, this algorithm is not applicable for all functions, therefore an underlying understanding of derivatives and calculus is required to determine applicability.

One of the key breakthroughs in this domain was backpropagation. During training, after each forward pass through the neural network, backpropagation performs a backward pass and adjusts the model’s parameters (weights and biases) with the aim of minimising the cost function (Kostadinov, 2019). The cost function is the measure of error between the value the model predicts and what the actual value is (Mulla, 2020). Backpropagation gained prominence through a seminal paper by Rumelhart, Hinton and Williams (1986). This paper outlined several neural networks where backpropagation worked faster than earlier approaches to learning, thus enabling neural nets to be used to solve new problems (Nielsen, 2015) contributing to the dominance of Deep Learning within the AI domain today.

2.2.6 The Future of AI and the Fourth Industrial Revolution

As the popularity and uptake of AI technology increases there are a number of concerns relating to safety, unfairness and inequality surrounding the use of such algorithms within our day-to-day lives (Schwab, 2017). The prominence and convergence of technologies such as AI, Cloud Computing and concepts such as the Internet of Things has led to many experts concluding we are entering the fourth industrial revolution (Maynard, 2015). The term ‘Fourth Industrial Revolution’ was created by Klaus Schwab, founder and executive chairmen of the World Economic Forum. Schwab states the importance of constructing the fourth industrial revolution to “ensure that it is empowering and human-centred, rather than divisive and dehumanizing” (Schwab, 2017, p.4).

There are many concerns surrounding the fourth industrial revolution such as the potential for mass unemployment due to automation which may exacerbate current inequalities

(Peters, 2017) due to the dominance of men within Computer Science and Engineering roles. These skills will be in increased demand within the coming years (Schwab, 2017) where training in Computer Science will be seen as a form of “4IR literacy” (Penprase, 2018). One of the ideas, identified by policy makers to mitigate against the rise in unemployment is to promote and situate lifelong learning as a way to equip individuals to deal with the rise of AI and to give individuals opportunities to meet these new workplace demands (Eynon and Young, 2020). The HE sector will play a key role in helping society to transition to the fourth industrial revolution (Gleason, 2018a), not only through subject specific education, but also via helping develop “soft skills” which will be imperative within the modified workplace. The World Economic Forum reported in 2016 in their Future of Jobs report (World Economic Forum, 2016) some of the top skills required by employers in 2020 including complex problem solving, critical thinking, emotional intelligence, judgement and decision making and cognitive flexibility. Alongside these skills, people will need to be able to frequently update their knowledge and learn new skills and connect the different aspects of their learning (Lewis, 2018) as well as be adaptable and flexible to keep up to date with the changes within the workplace (Gleason, 2018b). Therefore, AI education within the HE sector will play a key role within the fourth industrial revolution to ensure individuals have the technical and interpersonal skills required for the changing workplace. Identifying a framework of best practice will enable insight into the most appropriate pedagogies to teach these skills.

Concerns were raised regarding who is involved in the development of these applications and whether these developers were trustworthy and had good intentions (Ipsos MORI, 2017; Centre for Data Ethics and Innovation, 2020). Greater emphasis on data ethics and further transparency in the creation of these algorithms is necessary to continue the current level of engagement within both industry and the wider general public (Ipsos MORI, 2017) and to continue building trust in these technologies (European Commission, 2020). When determining current AI education provision it is pertinent to comprehend the extent to which the ethical issues are currently taught and highlighted to students.

2.3 Data Ethics

The rise and availability of so called “big data” has issued in a societal change where various human pursuits and decisions are being shaped by predictions created through the use of data. These predictions are used to influence activities such as shopping, medicine, law

enforcement and education (Richards and King, 2014). Ethical concerns arise relating to big data in regards to “the nature of what is being processed and who the processing is being done for or by” (O’Leary, 2016). The ubiquity of personal information being collected and used by corporations and governments such as location history, social network connections, search history and facial recognition are covered by legal and commercial secrecy which causes a lack of transparency (Richards and King, 2014). It is important to ensure that new practitioners within this domain are educated in the risks and ethical implications of this form of technology.

There are a number of potential hazards associated with Artificial Intelligence centred around inaccessible decision making, injustice including discrimination based on gender, racial or ethnic origin, disability or sexual orientation as well as growing concern relating to this technology being used for criminal purposes (European Commission, 2020). There are a number of issues surrounding the use of data, specifically within AI applications, including one of the most prominent concerns, bias. Bias can occur when an unrepresentative dataset is used to train an AI system, in turn leading to a system which can lead to discriminatory practices (Centre for Data Ethics and Innovation, 2020). It is important to consider the origins of the data. The majority of data used within these systems originates from humans and is collected by humans which incorporates an element of subjectivity. This subjectivity can lead to human oversight and the use of this data without identifying and mitigating against potential bias (Select Committee on Artificial Intelligence, 2018; European Commission, 2020).

A comprehensive understanding of the data to be used within any AI system is a necessity as well as understanding the potential issues which may arise from the combination of that dataset being used within a specific model. A dataset which inaccurately represents society, or even a dataset which accurately displays the unfair aspects of society all induce bias (Select Committee on Artificial Intelligence, 2018). These inaccuracies can be difficult to detect when the teams working on these system are comprised of people from dominant groups, their specific perspective dominates the decisions being made to the detriment of other identities and perspectives (D’Ignazio and Klein, 2020). Therefore, widening participation in educational sectors for AI courses should be a priority as well as reskilling individuals to ensure a wider scope of the population are data literate.

As well as the potential for AI technologies to be exploited for criminal purposes, there is also growing concern surrounding the rise in misinformation, relating to the actual technology and its potential impact on society as well as the technology itself being used to manipulate and disseminate false and inflammatory material (Centre for Data Ethics and Innovation, 2020). Russell (2019) argues that mental security and the “right to live in a largely true information environment” is currently under attack and that we are incredibly exposed to the “technology of misinformation.” The prevalence of misinformation could lead to deterioration in public trust in AI and technologies which centre around data analysis as well as affect individual autonomy surrounding consequences of customised content (Centre for Data Ethics and Innovation, 2020). Therefore, upskilling individuals and increasing the level of AI literacy may mitigate against the rise in misinformation in that people will be at a greater advantage to spot the false material.

Although there has recently been an uptake in the conversation regarding data ethics, as a consequence of the more high-profile cases of misuse and discrimination, there is still a lack of focus on ethical practices and consideration of risk within large companies. According to the 2019 Artificial Intelligence Index Report (Perrault *et al.*, 2019, p.6), “Only 19% of large companies surveyed say their organizations are taking steps to mitigate risks associated with explainability of their algorithms, and 13% are mitigating risks to equity and fairness, such as algorithmic bias and discrimination.” This scarcity of consideration is alarming when acknowledging the consequences and risk associated with deployment of these types of algorithms. Therefore, it is of pressing importance that courses within AI instruct on the issues and inform on strategies to mitigate these risks.

Current areas of interest and frequently mentioned areas related to the ethical challenges of Artificial Intelligence, particularly Deep Learning algorithms are fairness, interpretability and explainability (Perrault *et al.*, 2019). Interpretability and explainability are often used as synonymous terms. However, interpretability refers to the ability to observe cause and effect situations or conditions within the system. Whereas explainability refers to the explanation of the features of the interpretable domain which have contributed to produce a decision (Nassih and Berrado, 2020). It is of importance that AI practitioners have an understanding of these terms to ensure that they are cognisant of the many ethical challenges of AI. Transparency and reflexivity are also important considerations as they permit the people

involved in AI projects to clearly communicate their methodology (D'Ignazio and Klein, 2020). Transparency is particularly important when the system is being created for use with consumers, however, the level of transparency and explainability will be different depending on who is looking at it, whether this is developers, users, investigators or regulators (Select Committee on Artificial Intelligence, 2018).

Overall, there is a crucial and timely need for developers and researchers within this field to have greater comprehension of data ethics, to not only be aware of the consequences but to be able to bake in ethical practices within every stage of development. There is also a need to create development teams which are representative of society to address the biases inherent within data and within developers themselves (Select Committee on Artificial Intelligence, 2018). Training will be an important process to ensure that graduates are leaving courses with a strict understanding of ethical principles and codes of conduct and how to incorporate this within the development cycle.

2.3.1 Teaching Data Ethics

There are presently two main approaches to teaching data ethics; in individual courses which focus solely on ethics and policy, and integration of ethical principles into modules which make up the computing curriculum (Perrault *et al.*, 2019). Penprase (2018) advises that the “curriculum needs to help students develop the capacity for ethical reasoning, for awareness of societal and human rights” and to be able to assess the effects of the fourth industrial revolution technologies on individuals. Ethics, although included in the overall educational outlook are often not taught as a core aspect of Machine Learning courses (Saltz *et al.*, 2019). One of the main issues when ethics is taught as part of a Machine Learning module is that the method of integration into the content suggests that this topic is supplementary and not a core aspect which should be incorporated into all of the elements of the work (Saltz *et al.*, 2019). However, experiences teaching specific data ethics modules has proven popular with students and have had strong levels of engagement (Henderson, 2019). Therefore, a review of the current ethics provision within this study will identify how widely this topic is currently offered.

2.4 Representation and Diversity

As mentioned in Section 2.3, diversity and representation within the teams developing AI systems is an important factor in making sure these technologies are fair and do not exacerbate inequalities. This is important to ensure that the decisions being made do not exclude other “identities and perspectives” (D’Ignazio and Klein, 2020). It is postulated that “increasing the diversity of the workforce developing AI systems will reduce the risk that they generate discriminatory and unfair outcomes, thus ensuring that their benefits are more widely shared” (Stathoulopoulos and Mateos-Garcia, 2019). However, a major barrier to improving diversity within this domain is the current shortage of available data and statistics on diversity in industry and academia (Perrault *et al.*, 2019; Stanford University, 2021). This is especially pertinent relating to intersectionality, including sources of inequality such as class, race, ethnicity, religion, sexual orientation and age which need to be included in any analysis (Young, Wajcman and Sprejer, 2021). Young *et al* (2021) also recommend that large tech companies should be subject to reporting requirements relating to the “gender composition of their data science and AI teams”.

From the minimal data available, 78% of professionals globally within this domain are male, with only 22% female (World Economic Forum, 2018). This reduces to 20% women within the UK (Young, Wajcman and Sprejer, 2021). These findings only report on gender as a binary construct therefore these statistics do not represent a full picture of workforce gender demographics. An initiative funded by the Office for Students (OfS) UK to offer scholarships to underrepresented students within higher education to study a postgraduate conversion course in AI or Data Science has boosted the numbers of underrepresented groups, particularly of women enrolled on such courses as shown in Figure 2 (Office for Students, 2021a). Opportunities such as these are a key method for widening participation within this domain.

Representation within the academic pipeline is poor with the majority of AI department faculty identifying as male. On average 80% of AI professors are male (Perrault *et al.*, 2019). Increasing representation within academia may have a number of potential benefits including publication of AI research which is more “applied and socially aware” according to a report by Stathoulopoulos and Mateos-Garcia (2019). Out of the 25 most notable terms of academic papers co-authored by women, terms such as “fairness, human mobility...health” were the

top ones mentioned (Stathoulopoulos and Mateos-Garcia, 2019). Increasing representation within AI will not only decrease the current skills shortage identified but is also a necessity as outlined in Section 2.3 Data Ethics. A diverse cohort of students may also experience a variety of barriers to learning which is pertinent to RO1.

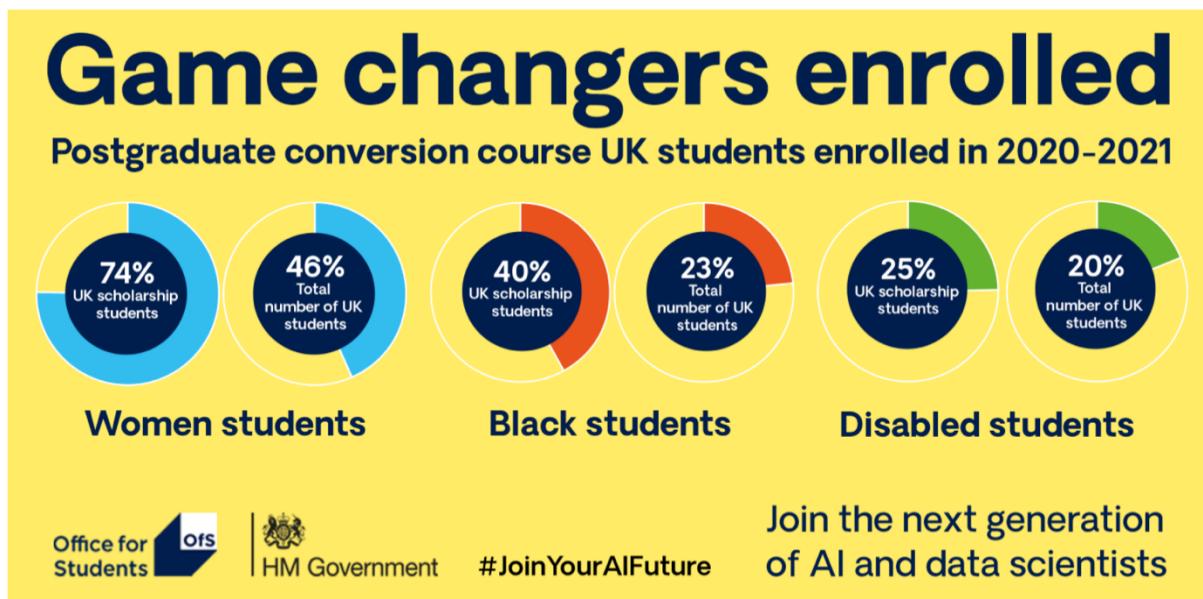


Figure 2: The percentage of underrepresented students enrolled on the postgraduate conversion course (Office for Students, 2021a).

Comparing representation in AI with other computing domains has shown that the issue of underrepresentation is prevalent throughout computing. Alongside the lack of diversity, professionals within this domain have often experienced discrimination. Within Cyber Security “41% of black cyber professionals feel they have experienced discrimination over their ethnicity” and “female respondents who felt they had experienced gender discrimination accounted for the largest proportion (23%) of all incidents” (KPMG and National Cyber Security Centre, 2020). Young, Wajcman and Sprejer, (2021) found that women working within AI and Data Science have a high turnover and a higher attrition rate than men. The findings from these studies indicate that a change is needed within the workplace culture within these domains to make these environments more inclusive.

2.4.1 Widening Participation

Widening participation, “the process of increasing diversity and representation of students who have traditionally not studied in higher education” (Jones and Thomas, 2005), is a concern for many universities and an issue many are endeavouring to improve upon within the immediate future (Rainford, 2020). One of the central tenets and objectives within the

English UK higher education sector is to “ensure that students, from all backgrounds (particularly the most disadvantaged), can access, succeed in, and progress from higher education” (Office for Students, 2021b). Widening participation is concerned with expanding educational opportunities to a diversity of individuals including people who are “economically, educationally and socially disadvantaged, in terms of poverty or social class, and also age, ethnicity or race and by gender” (David, 2010, p.3).

There are many different approaches to widening participation, however it is advised that “education should be changed to permit it to both gauge and meet the needs of underrepresented groups” (Jones and Thomas, 2005). Jones and Thomas (2005) also recommend that an institution’s actions should “be underpinned and informed by valuing and learning from difference and diversity.” The institutional approach to widening participation and the subsequent cohort diversity may impact differing barriers to learning students face and as a consequence their course satisfaction and attainment (RO2.e). As well as the fundamental change needed to ensure that universities are welcoming and catering to the needs of underrepresented cohorts of students there are also specific schemes which institutions can support which have shown success. For example, the Conference on Neural Information Processing Systems (NeurIPS) has co-located workshops from both the Black in AI and Women in Machine Learning groups which have both shown an increase in participation and submitted research papers since their inception (Stanford University, 2021). As discussed in the previous section representation with the overall aim of widening participation within AI is incredibly important to ensure a high level of AI literacy and to fill the current skills gap. Although not a direct research objective, these issues are pertinent to all of the research objectives to ensure educational strategies are appropriate and embed and encourage methods to encourage widening participation.

The growing use of blended learning and technology to assist and deliver teaching has enabled greater flexibility and democratization of learning (Yang and Cheng, 2018). This flexibility enables students to fit their course in around other work or caring responsibilities (Office for Students, 2021a) as they are able to complete assignments and coursework at any time and are not required to be on campus as much as traditionally situated courses (Poon, 2013). Consideration of this learning approach in attainment of RO2.b will be investigated.

There is a need for greater signposting of the opportunities for retraining within this domain as research conducted by Deloitte and Institute of Coding (2018) found that 54% of women working in non-digital areas were not aware of the opportunities for retraining or they did not believe they had the necessary skills to work in this field. Therefore, conversion courses and scholarship opportunities, such as the one offered by the Office for Students (2021a) are vital to widening participation within the field of AI.

2.5 Learning Theory

This section of the literature review discusses a number of aspects of learning and learning theory, from how we learn to the various educational learning theories and models and taxonomies. Comprehension of these aspects of learning are pertinent to the fulfilment of the research objectives. Being able to identify the learning theories employed within the modules under review (RO1) will ensure that commonalities are identified. Analysing the differing theories of learning will also highlight appropriate models and taxonomies for inclusion within the learning resource (RO2.b) and identify current educational best practice (RQ1).

2.5.1 The Concept of Learning

Learning can be identified as “any process that in living organisms leads to permanent capacity change and which is not solely due to biological maturation or ageing” (Illeris, 2007, p.18). There are a number of different interconnected conditions which can influence the learning process, Figure 3 displays these and the main aspects to learning. Illeris (2018) outlines in this framework that all learning integrates both an external interaction process which occurs between the learner and their “social, cultural or material environment” and an “internal psychological process” (Illeris, 2018, p.2). Both the internal and external conditions can impact upon the learning process and can themselves be influenced by a person’s disposition or the context of the learning space.

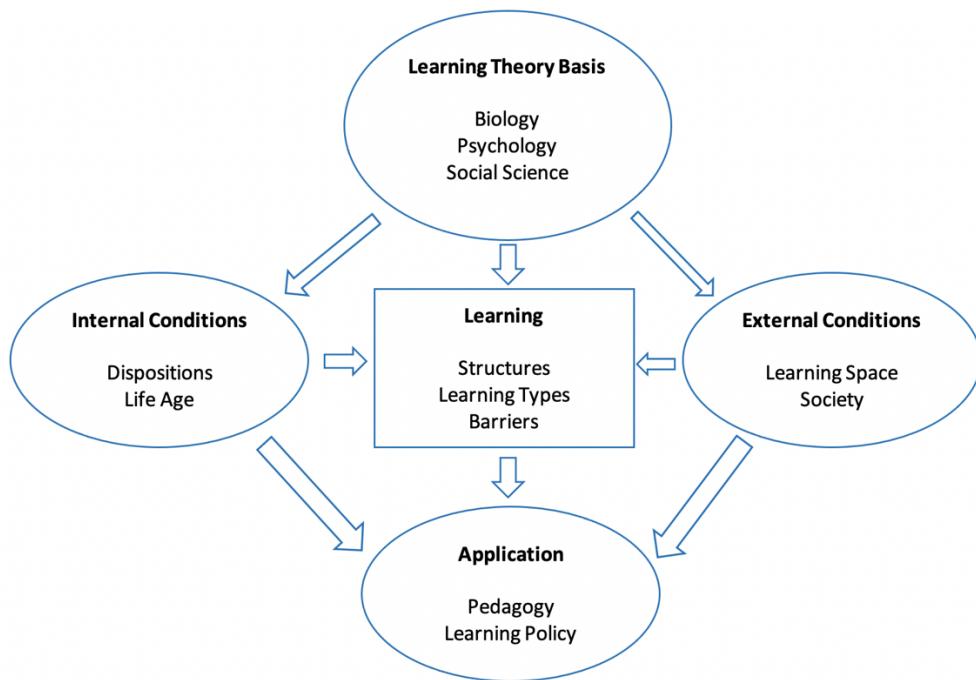


Figure 3: The main aspects of learning (based on (Illeris, 2018))

Each of the aspects of learning outlined in Figure 3 map to the research outcomes and questions defined in Chapter 1. For example, the Learning aspect pertains to the barriers students may encounter, this relates to RO1 and RQ2. Pedagogy is a central focus within the Application aspect which relates to RQ1 and RO2.e in determining the best practice for teaching this domain. The Learning Theory Basis can be mapped to RO2.c and RQ3 relating to cognitive barriers such as a lack of metacognition or self-efficacy. Internal Conditions relates to cohort diversity and educational background and the potential implications this may have on learning.

Demetriou and Spanoudis (2018, p.4) advise that it is a widely held belief that “the human mind (i) learns actively from interactions with the world, (ii) represents its interactions in mental models which are used to (iii) guide further encounters, especially when they are new and unexpected.” It is the response to unexpected situations and the ability to adapt which can provide a gauge of the differences of individuals and a reflection on intelligence (Demetriou and Spanoudis, 2018). Intelligence is associated with a number of sophisticated cognitive operations requiring comprehension of both our personal knowledge abilities and skills as well as those of others (Luckin, 2018). Therefore, our ability to accurately reflect and assess our knowledge and learning capabilities is inherent to the learning process.

2.5.2 Theories of Learning

The inception of learning theories can be traced back to philosophers such as Plato, whose epistemology was an “attempt to understand what it was to know” and Locke, who pursued an explanation of “the operations of human understanding” (Steup and Neta, 2020). Various theories discussed within this section have built upon these foundational ideas. Engeström (2018, p.46) defines four core questions which any learning theory must answer:

1. Who are the subjects of learning: how are they defined and located?
2. Why do they learn: what makes them make the effort?
3. What do they learn: what are the contents and outcomes of learning?
4. How do they learn: what are the key actions of processes of learning?

The consideration of these questions can assist when determining personal epistemology and aligning a learning theory to current teaching practice. Hattie and Donoghue (2018) advise that the science of learning is understanding which learning strategies are effective, however, understanding the context of when to use different strategies and with which student cohorts is the art of teaching. Determining the best practice for teaching the AI domain will require identification of the learning strategies currently employed and identifying whether any potential strategies would be more effective within this context.

The information processing theory has dominated the field of cognitive psychology since the early 1950s (Demetriou and Spanoudis, 2018). This theory posits that the human brain is similar to a computer in that “controlled attention, speed of processing, working memory, and inference are considered important in registering information, understanding, learning and, problem solving” (Demetriou and Spanoudis, 2018, p.7). This theory is intrinsically linked with the field of Artificial Intelligence, particularly early research within this field. Cognitive modelling and therefore the interdisciplinary field of cognitive science unifies computational AI models with experimental psychology techniques to create testable theories of the human mind (Russell and Norvig, 2013).

Relating to cognitive science, Piaget is the founding father of the cognitive development discipline within the study of intelligence (Demetriou and Spanoudis, 2018). Although Piaget worked primarily within child development, his theory of learning has relevance within adult

education and has been applied within the HE sector (Sutherland, 1999) and specifically within Computer Science education (Gluga *et al.*, 2013). Gluga *et al* used Neo-Piagetian theory, a derivative of the classical theory which can be applied irrespective of age, to assist Computer Science practitioners in evaluating the difficulty of programming learning tasks. Findings discussed how this enabled educators “to be more systematic in designing their learning and assessment materials” due to the mapping to Piaget’s cognitive development stages (Gluga *et al.*, 2013).

Two central components of Piaget’s position on learning are assimilation and accommodation. Assimilation involves the merging of new information with extant knowledge structures, however these structures are not changed (Reinking, Labbo and McKenna, 2000). Whereas with accommodation, new information requires the restructuring of existing knowledge to accord with this new knowledge, eventually transforming the learners views (Reinking, Labbo and McKenna, 2000). Piaget believed that learning is active and constructive, active in that the learner carries out a number of actions to build their mental models and understanding and constructive in that these actions coordinate into more progressive levels of functioning and understanding.

Piaget also proposed “reflective abstraction” a mental process which “amalgamates a series of actions and their results into a new structure” (Demetriou and Spanoudis, 2018, p.39). This new structure can help consolidate the students understanding and eliminate any confusion or inconsistencies prevalent within previous stages of learning. Essentially, reflective abstraction is the process of reflecting upon the thinking and learning process. The process of reflection is pertinent to the cognitive mitigation strategies proposed within this research (RQ3, RO2.c).

Piaget’s theory has been incredibly influential and has shaped a range of learning theories over the years, including one of the most dominant theories still widely used today – constructivism (Demetriou and Spanoudis, 2018). Constructivist theories centre around the ways in which learners create their own mental structures when interacting with the environment with a pedagogical focus on task-orientation and engagement in self-directed tasks (Wenger, 2018). The educational environment should facilitate a diverse range of interaction opportunities, including collaboration and should enable learners to enact

personal control over their learning (Bostock, 1998). This self-regulation of learning and reflection of appropriate learning strategies is emblematic of andragogy.

One of the core ideas relating to the constructivist viewpoint is the concept of active learning. Examples of active learning include the use of the students skills and knowledge to create something which physically embodies their knowledge or creation of solutions to problems posed to the students (Bostock, 1998). Relating to the problem posing educational strategies, to be effective, these scenarios must closely model the real-world issue being modelled and must be an authentic experience to provide a meaningful experience for the students. Problem-posing education is pertinent to Computer Science education due to the inherent need for Computer Scientists to have a level of problem-solving capability.

Experiential learning theory also uses the concept of active learning/experimentation and is influenced by the works of learning theorists Lewin, Dewey, Vygotsky and Piaget (Kolb, 2015; McCarthy, 2016). The Experiential Learning Theory (ELT) was created by David Kolb in 1984 and describes how experience translates into learning and knowledge (Kolb, 2015). ELT highlights the influential connection which can be developed within the classroom and the “real world” and emphasises the position of formal education in the promotion of lifelong learning (Kolb, 2015). This theory is to be applied as a process which “merges experience, perception, cognition and behavior” (McCarthy, 2016) and is therefore relevant to the research questions and objectives pertaining to the identification of difficulties (RQ1 and RQ2) and also cognitive mitigation strategies (RQ3 and RQ2.c). The four processes of the learning model are shown in Figure 4.

The four-stage cycle starts with Concrete Experience, these experiences provide a foundation from which to reflect on which entails the next stage of the cycle, Reflective Observation. The reflections and observations gathered from this second stage inform the Abstract Conceptualisation. The learner develops rules specific to the experience or applies existing theories relevant to it. The final stage in the cycle, Active Experimentation requires the learner to construct new ideas and methods to modify their next learning experience, which brings the cycle back round to Concrete Experience.

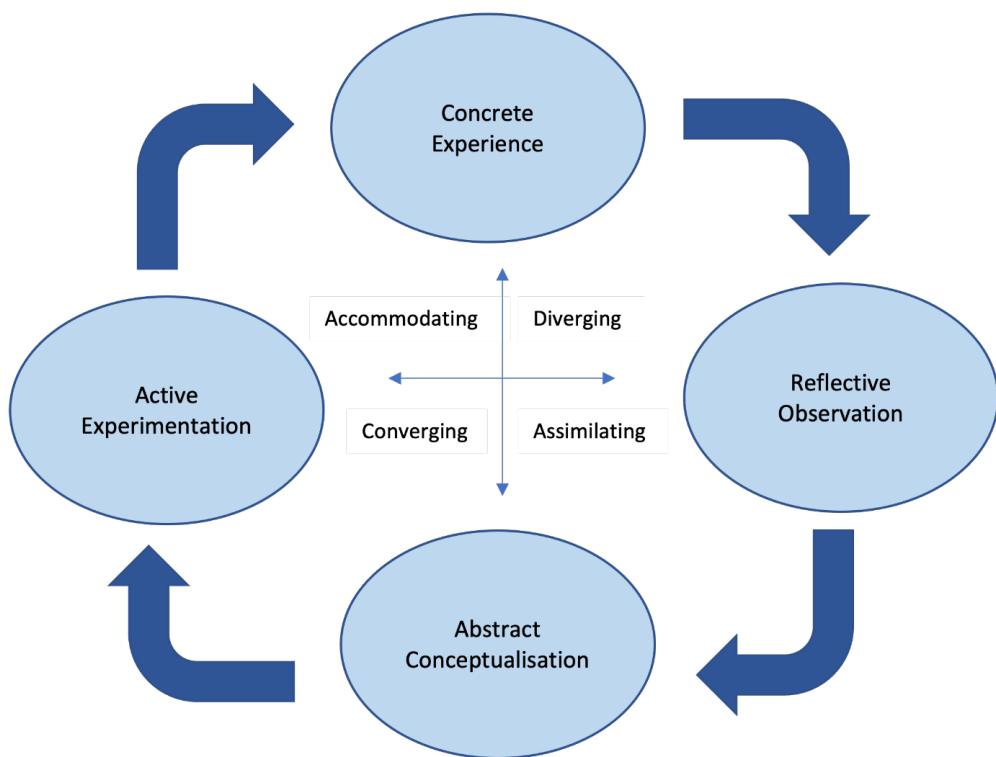


Figure 4: Kolb's Experiential Learning Cycle

Kolb also created the Learning Style Inventory which is based on the four-stage cycle. The use of the learning styles approach is widely applied, particularly in higher education, however there is no evidence to support teaching students in accordance to their 'learning style' as a successful education method, this disparity between evidence and practice has led to controversy and the branding of learning styles as a myth (Newton and Miah, 2017). There is instead a wider call to use pedagogical techniques which are demonstrably evidence based (Newton and Miah, 2017). Identifying evidence-based learning strategies relevant to the teaching of AI is a key aspect of this research and will be reviewed in the following section.

2.5.3 Learning Models and Strategies

Alongside theories pertaining to how learners acquire knowledge, there are also learning models/methods which provide specific strategies which the educator can implement within their practice to assist with student learning. Being cognisant of the differing learning models will enable identification of the use of these within the current AI education provision leading to awareness of what is currently considered best practice (RQ1). Investigation of these differing strategies will also allow consideration of how or if these are appropriate for facilitation within the learning resource (RQ2.b).

The Hattie and Donoghue (2016) model is composed of three different inputs and outcomes (skill, will and thrill) and three phases of learning (surface, deep and transfer). At the centre of the model is the idea that the student brings to the learning experience “a pre-existing set of personal qualities, abilities, knowledge and histories, all of which may impact their subsequent learning” (Hattie and Donoghue, 2018, p.99). Within Computer Science education this could pertain to educational background for example prior programming experience as well as cognitive barriers such as mathematics anxiety or low confidence in technical skills. The first of the three inputs and outcomes, ‘skill’, refers to the student’s knowledge and ability, for example their understanding of the domain from prior learning. This pertains to the student’s educational background and mental model of the domain. ‘Will’, refers to the mindset that students approach learning situations, this element of the model can be related to Dweck’s (2017) theory of the growth mindset which is situated in the belief that “everyone can change and grow through application and experience” (Dweck, 2017, p.7) irrespective of their aptitude or initial knowledge level. One of the most pertinent aspects to learning of the growth mindset is the enthusiasm for challenging yourself and persevering when encountering difficulties, this quality can be developed and potentially be an outcome of application of the Hattie and Donoghue model and is pertinent to RQ3 and RO2.c. The growth mindset and the ability to persevere in the face of difficulties is a key skill to assist learners through mastery of the domain threshold concepts (discussed in Section 2.7). Finally, the ‘thrill’ input of this model relates to the propensity the student has towards the subject being taught and their motivation to study this domain. Each of the inputs are also developed into outcomes, with each of the categories as valuable to the learning process and results.

One of the core ideas of the Hattie and Donoghue model is the success criteria and the importance of informing students of this early in the learning process. Awareness of the success criteria for a given task before commencing this work has shown to increase student’s goal-directed behaviour and be more strategic in their learning strategy selection, consequently they are more likely to enjoy the ‘thrill’ aspect of learning (Hattie and Donoghue, 2018).

Phases of learning are also outlined in this model including surface and educational deep learning. Surface learning occurs when the learner solely focuses and reproduces the main facts, it is also associated with anxiety related to learning and a disassociation with the

material (Draper and Waldman, 2013). Whereas, deep learning (in education as opposed to AI) enables learners to transform their knowledge beyond the main points and the reproduction of facts and critically engage with the material (Draper and Waldman, 2013). Both of these phases of learning are important, surface learning includes the acquiring of the subject matter vocabulary and content of the lessons through strategies such as note taking and imagery (Hattie and Donoghue, 2016). This surface learning then enters a consolidation phase where this knowledge is encoded so that it can be retrieved at a later date, the strategies for consolidation include practice testing, spaced versus massed practice and becoming competent at receiving feedback (Hattie and Donoghue, 2016). Consideration of how these strategies can be included in the learning resource (RO2.b) is pertinent to assist learners through the varying phases of learning.

Acquisition and consolidation phases are also part of the educational deep learning process, this phase consists of a series of learning strategies which promote the students abilities to reinforce deeper thinking and to enable them to be more strategic with their learning (Draper and Waldman, 2013). Strategies include self-monitoring and questioning, assistance from peers, reflection and evaluation and problem solving. When learners have at their disposal a variety of learning strategies, they become more self-regulated and tend to have high levels of metacognition (Hattie and Donoghue, 2016). Students may not like some of the phases of these types of learning, due to the necessity to practice alongside the “willingness to tolerate ambiguity and uncertainty during this investment phase” (Hattie and Donoghue, 2018, p.105) which in turn requires a commensurate level of metacognition. Therefore, improving student’s metacognition (RO2.c) is key to the attainment of educational deep learning. The Hattie and Donoghue (2018) model highlights the importance of students learning to learn. They suggest that learning skills such as critical thinking, resilience and the growth mindset are best developed relative to the specific module content and that incorporation of these strategies should be an integral part of the learning process.

This model also suggests that aiding the students to connect with their prior knowledge, boosting their confidence and reducing anxiety can make the most difference to the learning process and outcomes (Hattie and Donoghue, 2018). Determining potential sources of anxiety and barriers to learning is a key objective of this research (RO1) as well as strategies for alleviating these issues (RO2.c). The importance of student awareness of the criteria of success

is incredibly important to this model, this has shown that students will immerse themselves in challenging tasks if the learner feels the learning requirements are perceived as achievable with practice and feedback.

Another key learning model is problem-based learning (PBL). Research relating to the use of the problem-based learning model has shown that students learn a greater amount of content and have wider recollection when problem based learning is used as a mode of education with an authentic outcome (Gleason, 2018b). Determining the extent to which this pedagogical strategy is used within AI modules will enable identification of its applicability within this domain (RQ1). The PBL model is student centred and often involves group work to solve a variety of challenges, it has demonstrated its ability to motivate students and encourage greater engagement with the discipline content (Gleason, 2018b). As this model is student centred, the focus is to empower the learner so that they take responsibility for their own learning (O'Grady, 2012), enabling them to become more self-regulated. However, it is important to consider at what point in the learning to introduce PBL as it requires a sufficient level of surface knowledge to be able to solve the problem posed to them, students cannot move straight to deep and higher level thinking required for problem solving without copious content knowledge (Hattie and Donoghue, 2018).

Although it is possible to examine evidence from prior research to determine effective learning strategies and models which will be completed to determine current use in Computer Science education, determining the optimal strategy to use depends on where in the learning cycle the student is located (Hattie and Donoghue, 2018). The use of instructional design principles can assist in identifying the needs of the students.

2.5.4 Instructional Design Frameworks

Instructional design relates to the principles and processes by which learning materials, lessons and whole systems can be developed in an effective, reliable and consistent manner drawing from domains such as educational psychology and cognitive science (Molenda, Reigeluth and Nelson, 1983). The frameworks outlined within this section will be evaluated for applicability when designing the learning resource for RO2.b. One of the earliest theories of instruction was created by Gagné and Briggs in the 1960s which provided a comprehensive

framework within which most of the subsequent work in this area has been influenced by (Petry, Mouton and Reigeluth, 1987).

Bloom's taxonomy was created by Benjamin Bloom in 1956 in collaboration with colleagues, in which they published the framework for categorising educational goals (Armstrong, 2010). These educational objectives are grouped into a hierarchy containing six levels, at the lowest level is *knowledge*, with *evaluation* at the highest level and *comprehension* (level 2), *application* (level 3), *analysis* (level 4) and *synthesis* (level 5) in between (Moseley *et al.*, 2005). The creators of the framework indicated that all cognitive educational objectives could be identified within this hierarchy, however when educators utilise the hierarchy for their practice, there has been some confusion relating to where to locate particular educational objectives (Moseley *et al.*, 2005).

In 2001, Anderson and Krathwohl created a revision of Bloom's taxonomy to aim renewed attention on the taxonomy and to update it to incorporate the progress in our understanding of knowledge since the original publication (Moseley *et al.*, 2005). The revision maintains six cognitive process categories, however these categories were changed to: *remember*, *understand*, *apply*, *analyse*, *evaluate* and *create*. These cognitive process categories are ordered relating to the scale of their complexity. Figure 5 demonstrates how these new categories map to Bloom's original taxonomy. One of the major changes with the revision is the significance assigned to the use of the taxonomy in the planning stage of course creation, instruction and assessment and the importance of aligning the three (Moseley *et al.*, 2005).

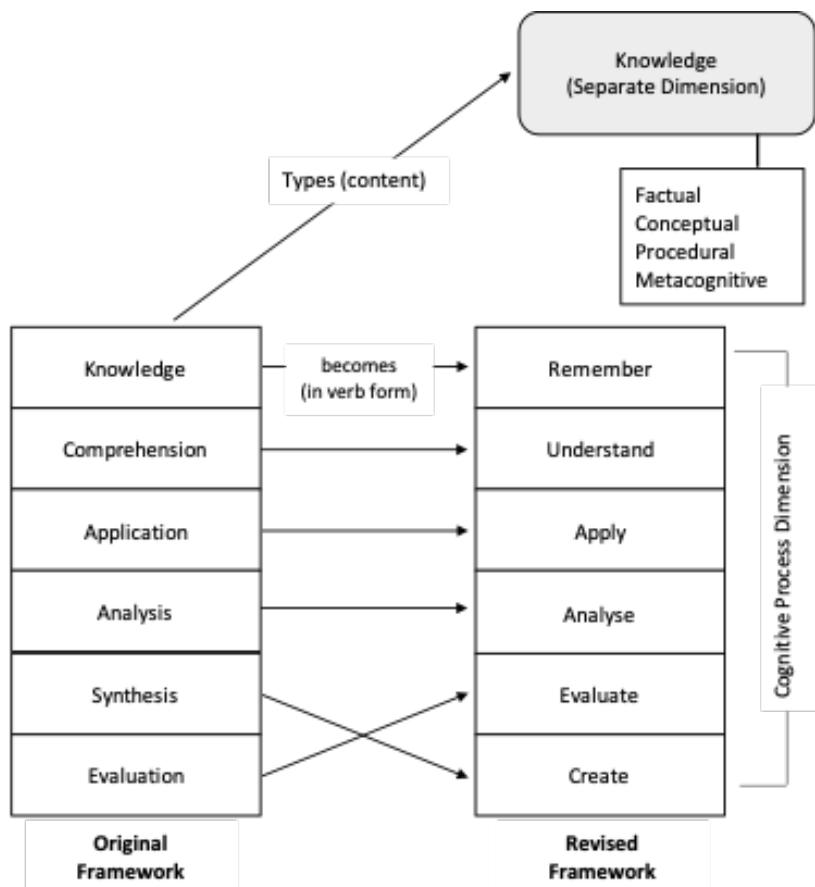


Figure 5: Changes to the Bloom taxonomy from Krathwohl and Anderson (image altered from (Moseley et al., 2005)

Anderson and Krathwohl's revision also changed the taxonomy from one dimension to two, where we now have the cognitive process dimension and the knowledge dimension. The knowledge dimension consists of four main categories, three of the categories maintained the main ideas/themes of the original taxonomy, however the fourth category – metacognitive knowledge is new (Krathwohl, 2002). The four knowledge types are:

- Factual Knowledge** – “the basic elements that students must know to be acquainted with a discipline or solve problems in it.” For example, understanding subject specific terminology.
- Conceptual Knowledge** – “the interrelationships among the basic elements within a larger structure that enable them to function together.” For example, comprehension of categories, principles, theories and models pertinent to the domain.
- Procedural Knowledge** – “how to do something; methods of inquiry, and criteria for using skills, algorithms, techniques, and methods.” For example, understanding of

domain specific skills, techniques and methods and the understanding of when is appropriate to use each of these.

- d) **Metacognitive Knowledge** – “knowledge of cognition in general as well as awareness and knowledge of one’s own cognition.” For example, understanding different cognitive tasks and self-knowledge.

(Source: Krathwohl, 2002)

The emphasis on the cognitive domain has led to some criticism of both the Bloom taxonomy and the revision by Krathwohl and Anderson, however this emphasis is pertinent to the aim of this research in determining some of the potential barriers to learning (RQ1) and the focus on cognitive mitigation strategies (RQ3, RO2.c). The revised taxonomy and its categorisations should be considered to assist in scaffolding the students learning when undertaking courses within the AI domain. Consideration of the four knowledge types outlined by Krathwohl (2002) will be key when creating the learning resource (RO2.b) and the creation of a framework of topics for a Machine Learning course (RO2.d) to ensure that all categories are investigated.

2.6 Computer Science Education

This section endeavours to uncover existing best practice within the computing domain, particularly specific aspects of this domain which pertain to Machine Learning such as programming skills. A review of current research specifically outlining pedagogy for teaching AI and Data Science has also been conducted to determine the prevalence of the learning theories, models and instructional design principles outlined in the previous section to identify current best practice (RQ1).

Identifying best practice and pedagogy for Computer Science education is a thriving research field in part due to the complexity of teaching such a technical subject to a diverse audience. As Drummond (2009) points out “Computer Science is one of the few sciences that do not necessarily specify an entry subject in their particular discipline” as a consequence students within computing enter the programme with diverse educational backgrounds, resulting in irregularities in students a priori knowledge. Crick (2017) completed a systematic review of the existing literature pertaining to computing education and found that comfort level in course material and mathematics background were found to have a positive impact on success. Determining the barriers students face when learning AI, particularly in relation to

educational background and mathematics anxiety is a key objective of this study (RO1). Analysis of current best practice within computing education will allow for consideration of mitigation strategies to alleviate the barriers identified.

Relating to learning theories and taxonomies, Piaget's theory of constructivist learning has been suggested as a suitable pedagogy for Information and Computing Sciences and particularly for programming modules (Crick, 2017). This includes the use of active learning through exploratory tasks and the aim of helping learners develop an appropriate model relating to how to investigate and develop an understanding of key concepts. A key aspect of Piaget's theory is assimilation, which involves the linking of new and extant information. The student's educational background is key to the assimilation process. If the student has no a priori knowledge of the domain they will not have a conceptual model to link the new information to. Piaget also discusses the importance of reflective abstraction, in which the student reflects on the learning process. Reflective abstraction is a key process relating to RO2.c and is relative to a number of categories of Bloom's taxonomy of the cognitive domain, which is also extensively used in the design and moderation of courses (Crick, 2017). Bloom's taxonomy has been used within Computer Science education to classify levels of learning objectives and assessments, however issues were identified in the difficulty of applying this taxonomy (Gluga *et al.*, 2013).

2.6.1 Programming

Most undergraduate Computer Science courses include modules that specify teaching the theory, the propositional knowledge and practical application of that theory to ensure procedural knowledge (Tendre and Apiola, 2013). Questioning and exploration of concepts through computational modelling and programming enables students to further their engagement with the theoretical concepts being taught (Curzon *et al.*, 2018). This is a prevalent pedagogical approach based on constructivist ideas and active learning (Hazzan, Lapidot and Ragonis, 2011). As AI/Machine Learning modules are mainly taught within the computing domain, it is expected that these courses will follow this practice. However, the potential lack of prior mathematics knowledge needed to comprehend the theoretical aspects of AI may be exacerbated by the use of programming exercises as learners may have difficulty identifying suitable algorithms for a specific task or may not understand the architecture of a model and how to program this which in turn may affect their confidence and self-efficacy.

Therefore, it is important for the module instructor to identify the cohort level of mathematics knowledge at the outset of the course. Programming modules have a number of similarities with AI, including the variation of abilities and educational backgrounds of the cohort, which can lead to a diverse range of successful and unsuccessful outcomes (Robins, Rountree and Rountree, 2003).

There are a number of potential obstacles which can hinder learners progressing from novice to expert within programming including confidence and lack of strategies and mental models (Robins, Rountree and Rountree, 2003). Mental models are “representations of situations with a strong iconic component somehow depicting the represented situation to the thinker” (Demetriou and Spanoudis, 2018, p.14). The obstacles can manifest when students have a lack of social persuasion, such as feedback from peers (Loo and Choy, 2013) which can boost self-efficacy (Ramalingam, LaBelle and Wiedenbeck, 2004), as well as a deficiency in “soft skills” such as goal setting and critical thinking (Grow, 1991). The potential obstacles encountered by novice programmers are relevant to a number of aspects of AI courses due to the preconditioned ideas students may have regarding the domain and computing and mathematics understanding from previous courses. For example, logically verifiable algorithms have been central to the theory of computing practice, however this differs with Machine Learning as a typical model is likely to be abstruse and the verification process “is not a logical proof of correctness” (Shapiro, Fiebrink and Norvig, 2018). There are also concerns relating to the fact that learners are often taught how to write programs, but not necessarily shown how to read them, which can have implications on the students code literacy (Crick, 2017). This disparity in code literacy can have wider implications for students when required to apply their programming skills within a different domain, such as Machine Learning.

A key strategy which has been employed to assist students with their progression to expert includes the use of mental models (Ramalingam, LaBelle and Wiedenbeck, 2004). A mental model invokes a user’s internal representation of components and rules of a system. There may be variation regarding the completeness of the mental model as a synonym for the learner’s comprehension of the modelled domain (Canas, Bajo and Gonzalvo, 1994). A student’s mental model can be embellished through experiential learning tasks (Ramalingam, LaBelle and Wiedenbeck, 2004). Research has demonstrated that a learner’s mental model

influences their self-efficacy and this consequently can affect course performance (Ramalingam, LaBelle and Wiedenbeck, 2004). Examining student's preconceptions of Machine Learning will enable a clearer understanding of their mental model prior to a module within this domain and any misconceptions they may have.

2.6.2 AI, Machine Learning and Data Science

There is currently a lack of research pertaining to the best practice for teaching Machine Learning to a HE audience. This is a concern as the number of students who enrolled or applied to partake in an introduction to AI or introduction to Machine Learning course has increased by almost 60% within the past four academic years (Stanford University, 2021). The majority of AI/Machine Learning courses are taught at masters level (Perrault *et al.*, 2019). University students are a key source of talent for companies seeking to recruit for employees with data skills and these companies prefer to recruit new graduates from quantitative disciplines such as Computer Science and Engineering (Department for Digital Media and Sport, 2021). Therefore, it is key to ensure that higher education institutes are equipping students with the adequate skills and knowledge to fulfil these roles. A recent report quantifying the UK data skills gap (Department for Digital Media and Sport, 2021) identified the top ten skills that businesses say graduates are lacking (Table 1). Included within these skills were Machine Learning, data processing, data ethics and programming, all skills which should be included within a Machine Learning module.

Top 10 Skills Businesses Advise Graduates are Lacking	
Basic IT skills	Data Ethics
Machine Learning	Programming
Data Processing	Data Communication Skills
Knowledge of emerging technologies and solutions	Advanced Statistics
Analysis Skills	Data Visualisation

Table 1: Skills that graduates are lacking (from GOV.UK (2021))

Although there is a lack of research specific to teaching Machine Learning at HE level, there is a growing body of research relating to teaching this domain to school age children. Vartiainen *et al* (2020) recognise the importance of teaching children AI due to the ubiquity of this technology and that children are now growing up with this technology. They also press the importance that due to the potential of Machine Learning, both "beneficial and malign" there is a greater need for children to understand the "ML-rich technological world" (Vartiainen,

Tedre and Valtonen, 2020). Involving children with the exploration of Machine Learning systems has often not been investigated due to the “inherent difficulties in bringing such abstract and highly complex phenomena into young children’s creative grasp (Vartiainen, Tedre and Valtonen, 2020). Although there are identified difficulties, Vartiainen and colleagues carried out research to determine if children could grasp the fundamentals of Machine Learning. They used Google’s Teachable Machine (Google, 2021) with children between the ages of 3-9 years old and got them to train a Machine Learning model using a set of three facial expressions chosen by the children themselves. The children could explore the input-output relationships with the Teachable Machine with guidance and support from the supervisory adult who then interviewed the children to explore what they thought had happened and why, when they taught the computer.

The results from this study demonstrated that the children recognised the process of “being a teacher of a computer” and that this was new regarding their previous experiences (Vartiainen, Tedre and Valtonen, 2020). The adoption of simple easy to use Machine Learning tools provided the children plenty of room to explore and familiarise themselves with the AI domain without having to have programming experience or write syntax. The children also had little difficulty in instructing their peers on how to use the Machine Learning tool and displayed detailed reasoning on how to create an effective dataset and to evaluate if the system was classifying the inputs accurately (Vartiainen, Tedre and Valtonen, 2020). Although this study was undertaken with children, the results can be interpreted as potential ideas for teaching adults this technology. Vartiainen *et al*’s study identifies the benefits of allowing learners to explore the overarching concept of Machine Learning through exploration of easy-to-use tools before delving into the syntax to build a clearer mental model of this domain, this may be a potential idea to explore with students in HE. The study by Vartiainen *et al* also highlighted the importance of building and using an input dataset that the learners could easily relate back to themselves (such as facial expressions in the study) so that they could recognise the system as learning.

Code.org (2021b) a non-profit organisation aimed at widening access to Computer Science in schools and for underrepresented groups has recently expanded its offering to include an AI and Machine Learning module (Code.org, 2021a). The module consists of twenty-one lessons, including the lesson plan and resources for teachers culminating in the students producing a

Machine learning app based upon a variety of datasets which may be appealing to the student including “Nutrition” “Top Songs” and “Movie Stats”. Content on the module includes differing types of Machine Learning, classification models and an AI code of ethics. The content of this module is similar to modules within this domain offered at HE level, although the content of the material is at a much lower detailed level, the use of activities, instructional videos and practical activities to convey the material are relevant instructional tools at any educational level.

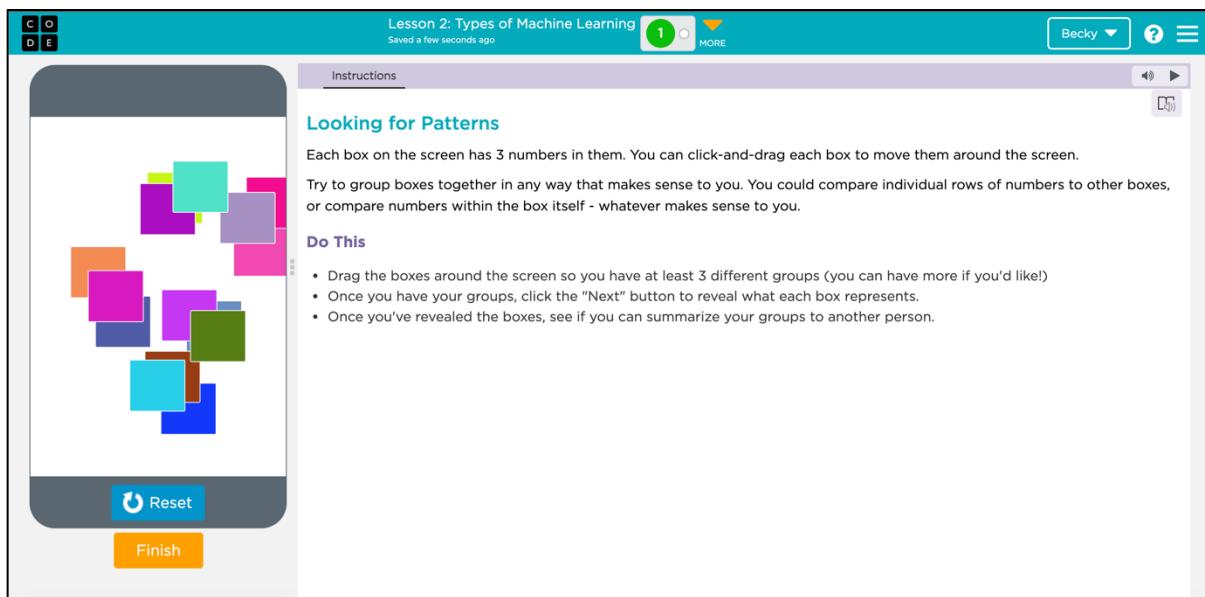


Figure 6: An example of one of the tutorials from (Code.org, 2021a)

As shown in Figure 6, the content of the tutorials is displayed in both a textual form and a visualisation to assist the learner in comprehending the material. Presenting the content in a variety of ways is pertinent to RO2.b and further exploration of these techniques will be relevant to the creation of the learning resource.

Due to the lack of research specifically focused on Machine Learning education, the literature search was expanded to include AI courses generally and Data Science. The underpinning concepts of these domains are markedly different from “traditional” computational curricula which usually advance from basic linear algorithms to control structures and data types then move on to include programming syntax (Selby, Dorling and Woppard, 2015; Vartiainen, Tedre and Valtonen, 2020). However, concepts within AI are noticeably different including the “shift from rule-driven to data-driven thinking, from transparent and explicit to opaque, from deductive to inductive, and from sensitivity to syntax to brittleness of models and sensitivity

to bias in training data (Vartiainen, Tedre and Valtonen, 2020). Therefore, there is a need for scaffolding of learner's mental models to accommodate this new way of computational thinking. Scaffolding techniques may be supportive (i.e. conceptual) or reflective (e.g. metacognitive) (Doo, Bonk and Heo, 2020) and include activities which build in complexity and provide an opportunity for feedback to allow the learner to incrementally build their mental model.

AI and Data Science courses often have similar content, relating to data processing (Brunner and Kim, 2016), although there is an element of ambiguity related to the term Data Science and the technical skills that are required. For example, companies looking to recruit a data scientist will often expect the candidate to be competent in AI (Asamoah, Doran and Schiller, 2015). Many Data Science courses contain principles related to Machine Learning (Hardin *et al.*, 2015) and may even cover this topic as part of the module content. From a technical perspective, data scientists "are seen as experts in advanced computational tools, data mining algorithms, statistical analysis, and machine learning" (Asamoah, Doran and Schiller, 2015). Similar to Machine Learning education, Data Science education is still a relatively young discipline (Schwab-McCoy, Baker and Gasper, 2021), therefore the existing body of research is limited compared to other strands of computing education. Although Data Science is often incorporated into statistics courses, as a result some elements of statistics education research were also reviewed.

Existing research relating to the teaching of Data Science has demonstrated that some courses within this domain do not require students to undertake a course pertaining to probability theory, multivariate calculus, linear algebra or statistics (Asamoah, Doran and Schiller, 2015), topics which are characteristically associated with introductory AI courses. Consequently, learners may have difficulty comprehending the heavy mathematical theory involved in such courses. Any difficulty pertaining to mathematics will be explored in relation to RO1. Recommended best practice for teaching a Data Science module involves the inclusion and embedding of a range of case studies, live code demonstrations within the lectures and minimisation of the use of mathematical notation and instead, where applicable, using computational approaches such as visualisation (Hicks and Irizarry, 2017). Determining the extent to which these strategies are currently employed will be investigated as part of the fulfilment of RO1 (p.6).

A study by Schwab-McCoy, Baker and Gasper (2021) into Data Science education sent out a survey to lecturers within this domain and asked what is the most challenging aspect of teaching this topic? The most frequent response was the scope of the course and the difficulty in narrowing down the vast amount of content into an introductory course, leading some respondents to state that they didn't have enough time to cover all of the topics they wished to. This finding is pertinent to RO2.d (p.6). The creation of a framework of topics should alleviate the challenges relating to the wealth of content. The instructors also identified the variety of learners educational background as a challenge with some students unprepared for the mathematics, statistics or programming aspects of the course (Schwab-McCoy, Baker and Gasper, 2021). Survey respondents also raised the requirement for more teaching resources, in particular interactive online tools which would enable them to incorporate active learning into their classroom. Online learning will be explored in Section 2.9 to determine the applicability of this as a method for hosting the learning resource in fulfilment of RO2.b (p.6). Pertinent to the identification of the best practice for teaching AI is the concept of the pedagogical content knowledge discussed within the following section.

2.7 Pedagogical Content Knowledge and Threshold Concepts

Upon analysis of existing pedagogical research, Lee Shulman, an educational psychologist identified a lack of research focused on the subject matter of the content of the course they were teaching. Very few of the researchers had enquired or discussed "how the subject matter was transformed from the knowledge of the teacher into the content of instruction" or how "particular formulations of that content related to what students came to know or misconstrue" (Shulman, 2013). Shulman believes that to optimally acquire a teachers full capability, equal attention should be paid to the content elements of teaching as has been dedicated to the teaching process (Shulman, 2013). To establish this theory of teaching, Shulman suggested a new kind of content knowledge – Pedagogical Content Knowledge (PCK), which extends the knowledge of the subject domain to encompass the understanding of the subject matter knowledge specific to teaching.

At the core of pedagogical content knowledge is the idea that it comprises of methods of "representing and formulating the subject that makes it comprehensible to others" (Shulman, 2013). This idea encompasses the most commonly taught topics within the domain, the most effective description of these ideas, and the use of appropriate demonstrations, visualisation,

analogies and examples. A key component of PCK is the awareness of issues relating to the content domain, such as difficult topics and which specific aspects make it difficult, what preconceptions the student may have of the content, which can be influenced by the varying backgrounds of the students. The concept of preconceptions and educational background are especially pertinent as the student brings this to their learning, which can influence motivation (Zepke, 2013). It is the role of the educator to restructure the understanding of the learners to eliminate any misconceptions or inaccuracies.

There are a range of methods which can be used to determine a subject PCK, for example within the field of Computing, interviews, questionnaires and observations were used by scholars to comprehend computing PCK (Hubbard, 2018). Relating to Anderson and Krathwohl's revision of Bloom's taxonomy, PCK can encompass many differing forms of knowledge including factual, conceptual, procedural and metacognitive knowledge (Krathwohl, 2002; Hubbard, 2018). Therefore, it is important to consider how teaching practice can incorporate support and inclusion for each of these types of knowledge.

Within a subjects pedagogical content knowledge, we also have the idea of threshold concepts. Zepke (2013) argues that threshold concepts are a feature of PCK as they provide a perspective of the course content that can have transformational and knowledge integrating benefits. Although some of the threshold concepts may be troublesome, identification of these can provide teachers as well as the learners methods of understanding this domain that would be unavailable without knowledge of these concepts.

Threshold concepts can be determined through a set of characteristics which identify them as core concepts which can potentially impede learning. These characteristics can be categorised as transformative, irreversible, integrative, bounded and potentially troublesome (Meyer and Land, 2003). The main characteristic of a threshold concept is its transformative nature, ensuring that a previously inaccessible way of thinking about a specific topic occurs. The irreversible nature of threshold concepts ensures that once a topic is understood it cannot be unlearned. Comprehension of a topic considered a threshold concept can reveal relationships between subject areas which were previously thought of as disparate, however threshold concepts can also be bounded, meaning that each concept does not generally explain the complete discipline. The basis of a threshold concept is that it is inherently troublesome to

learn, sometimes the concept may be counter-intuitive to beliefs a student may already have relating to a specific subject (Baillie, Bowden and Meyer, 2013). A study by Shinners-Kennedy and Fincher (2013) investigated the transformative and integrative nature of threshold concepts within Computer Science, they asked students to write a “transformative biography” in which the student identifies a computing concept which changed the way they think about this domain. Relating to the integrative aspect of threshold concepts they asked students to create a concept map relating to object-oriented programming. The purpose of these two tasks was to attempt to identify the threshold concepts within this domain, however the study found that insight from teacher’s PCK was a more fruitful approach to determining the threshold concepts.

Introducing students to threshold concepts indicates that teachers are keen for learners to be able to construct meanings for themselves by exploring strategies such as reflection and questioning, enabling them to evaluate their current knowledge and make connections with the threshold concepts (Zepke, 2013). However, Walker (2013) advises that threshold concepts are usually the aspects of the module where students ‘get stuck’ as they can be regarded as a “particular state of expert knowledge.” As a consequence, students who have not yet coherently understood a threshold concept may learn new ideas in a more disjointed manner as they cannot yet integrate this new concept into their way of thinking. However, students with advanced metacognition will be equipped to navigate through these threshold concepts as metacognitive processes are “associated with enhanced cognitive performance” (Luckin, 2018, p.45).

Once the threshold concept has been understood, the learner can then integrate different aspects of the subject into their analysis of problems (Land *et al.*, 2005). Preparation for threshold concepts within lessons should enable lecturers to implement strategies to assist the students when they encounter these concepts, demonstrating that they can countenance learner confusion (Cousin, 2006) and help them to formulate differing approaches to improve student understanding. Students can often have a complex and varying path to learning which unless communicated or identified by the lecturer, can often lead to miscommunication and confusion regarding student progress (Lucas and Mladenovic, 2007).

Identification of the threshold concepts within the domain of AI (RO2.a), specifically Machine Learning will enable better teacher understanding of the specific topics which can cause students difficulty and help inform and guide future learning design and best practice. Opening up the dialogue relating to threshold concepts within a specific domain enables the content experts to research and explore varying methods of assisting students to understand a difficult subject and facilitate discussion between the teachers and students aiming to expand the existing views of the subject matter and as a consequence integrate these new findings into the PCK (Zepke, 2013).

2.8 Cognitive Barriers

One of the objectives of this work as outlined in Chapter 1 (RO1), is to identify the potential barriers that may impact on student attainment in Machine Learning courses. As part of this exploration of potential difficulties, several cognitive barriers which can impact on and pertain to learning have been explored including mathematics anxiety, self-efficacy and confidence, with the aim of overcoming these difficulties to improve student metacognition and become more self-regulated learners.

Hattie and Donoghue (2016) identified that “negative emotions, such as those induced by fear, anxiety, and stress can directly and negatively affect learning and memory. Such negative emotions block learning”. Therefore, it is important to determine the prevalence of these emotions amongst students studying this domain and identify potential strategies to alleviate these barriers to learning (RQ3).

2.8.1 Mathematics Anxiety

AI, and subsequently Machine Learning, is an advanced subject which combines aspects from both mathematics and Computer Science where students may have a lack of confidence in their ability to do well in such courses, especially if they suffer from mathematics anxiety. Mathematics anxiety “involves feelings of tension and anxiety that interfere with the manipulation of numbers and the solving of mathematics problems” (Richardson and Suinn, 1972). Machine Learning tends to primarily be taught within Computer Science departments where students are taught programming languages and data analysis tools but rarely and in no great detail are taught mathematics and statistics (Deisenroth, Faisal and Ong, 2020).

As Machine Learning becomes more ubiquitous and the software packages become more higher level and easier to use, the lower-level technical details become more hidden away from the user. However, this poses a number of hazards as the practitioner will be “unaware of the design decisions and, hence, the limits of machine learning algorithms” (Deisenroth, Faisal and Ong, 2020, p.xi). Therefore, it is important, alongside integration of the ethical implications of this technology, to provide students with an understanding of the foundation of Machine Learning, including the mathematics and statistics. Having some understanding of the mathematical foundations of Machine Learning can enable greater comprehension of the fundamental principles upon which complex Machine Learning systems are built. Also leading to the creation of new Machine Learning solutions, comprehension and debugging of solutions and understanding the inherent assumption and limitations of these methodologies (Deisenroth, Faisal and Ong, 2020).

A recent review quantifying the UK’s data skills gap (Department for Digital Media and Sport, 2021) investigated the current curriculum content relating to mathematics and the level of mathematics literacy amongst the adult population in the UK (aged between 16-65) and found that a swift change was required to ensure that the population could have a chance at becoming data literate. The report also found that unlike higher performing peers, the UK does not place quantitative skills at the core of curriculum content. This is especially noticeable within tertiary education [education above school age] (Department for Digital Media and Sport, 2021) which presents a concerning challenge for modules within the curriculum which require a high level of mathematics knowledge, such as Machine Learning. A study by Hunt *et al* (2011) relating to mathematics anxiety levels within the British undergraduate student population indicated that “maths [sic] anxiety was significantly greater in women than men” and that there was a higher prevalence of mathematics anxiety within the science faculty than was anticipated (Hunt, Clark-Carter and Sheffield, 2011). This study consisted of 1153 participants (544 male, 609 female) from all university faculties who completed the Mathematics Anxiety Scale – UK (MAS-UK) (Hunt, Clark-Carter and Sheffield, 2011). These findings indicate that mathematics anxiety may be more prevalent within STEM cohorts and that this is more commonplace with female students. An analysis of the current AI/Machine Learning provision (RO1) will indicate the expectation level of mathematics skill for students within this domain as well as identifying specific mathematical concepts taught within such modules. Determining cohort demographics, mathematics attainment level and

confidence in mathematics skills of students studying within this domain should enable a characterisation to be built relating to student's mathematics profile and whether this topic is a potential barrier to learning (RQ2).

Mathematics anxiety can be associated with a number of negative consequences including avoidance of subjects which contain elements of mathematics (Ashcraft, Kirk and Hopko, 1998) and can affect the success of the most intelligent and determined students (Iossi, 2007). The variability of student cohorts in respect to societal and educational background has also been linked to educational motivation constructs including self-efficacy and anxiety related to mathematics (Lee, 2009). Mathematics self-efficacy and anxiety can be associated to the "notion of culturally diverse sources of self-formulation" (Lee, 2009). Subsequently, there is a correlation between anxiety and self-efficacy and the influence on academic performance, especially in mathematics (Ironsmith *et al.*, 2003). One of the potential explanations of this is that students who have a high level of self-efficacy generally display greater persistence and more sustained effort when encountering educational challenges (Parker *et al.*, 2014). Therefore, there is a focus in this research on cognitive strategies which may alleviate identified difficulties (RQ3).

Research relating to mathematics anxiety has identified a number of strategies to alleviate this issue. These interventions include methods aimed at affecting the whole cohort through curriculum changes and experimentation with differing psychological approaches (Hembree, 1990; Sarkar, Dowker and Kadosh, 2014; Ramirez, Shaw and Maloney, 2018). Psychological interventions in the past varied between behavioural strategies aimed at alleviating negative emotions towards mathematics and cognitive strategies which aimed to relieve concerns disclosed by the student (Hembree, 1990). The psychological interventions which proved the most effective included increased mathematics exposure, systematic desensitization, anxiety management and relaxation training (Hembree, 1990; Ramirez, Shaw and Maloney, 2018). Although these treatments displayed promise in alleviating this form of anxiety, these methods in many instances require specially trained staff who may not be readily available to carry these interventions out. Therefore, alternative methods to improve mathematics anxiety will offer greater opportunity to reduce this anxiety without the need for specialised staff.

One of the main benefits associated with the lowering of student level of mathematics anxiety is that it can increase mathematics assessment scores (Hembree, 1990). Ramirez *et al* (2018) suggest that insufficient self-regulatory processes in students can lead to lower perceptions of personal competence. Implementing strategies to help improve students' self-efficacy may ease mathematics anxiety and increase student confidence in their mathematics ability. This is something worth exploring in respect to research objectives RO1 and RO2.e (p.6).

Classroom interventions and curriculum strategies aimed at reducing mathematics anxiety include the use of games and interactive platforms, providing specialist equipment and using alternate techniques for course material delivery such as group work or tutorials (Hembree, 1990; Ramirez, Shaw and Maloney, 2018). Refreshing the format of lessons is an easily achievable objective as is the use of interactive platforms and peer interaction which may provide an attainable way to reduce mathematics anxiety within AI courses. This will be explored in relation to research objectives RO1, RO2.b and RO2.e (p.6).

A number of differing pedagogical approaches specifically relating to changing the student mindset have also been trialled such as changing the learners' perspective of failure as an important positive learning process and the viewing of mathematical problems as a challenge rather than a threat (Ramirez, Shaw and Maloney, 2018). In order to identify cognitive changes in students, lecturers need to be familiar with a series of learning interventions, such as identifying a student's prior understanding when they are approaching a task, having a comprehensive understanding of the material so that they can provide meaningful and challenging experiences to ensure progressive development and recognising when a student has achieved the learning outcomes (Hattie, 2012).

Other approaches such as retesting and self-paced learning have been shown to reduce learner's mathematics anxiety (Iossi, 2007). Retesting is a method used to motivate students to relearn skills and concepts which the learner had not yet initially fully comprehended to a mastery level (Juhler *et al.*, 1998). Self-paced learning can alleviate mathematics anxiety through alignment with learning goal orientation instead of performance goal orientation (Iossi, 2007). Another approach which has proven to be effective includes distance learning incorporating technology. This strategy has been deemed successful due to the "anonymity of an online course" (Iossi, 2007). Distance learning can be identified as the provision of

education at a geographically distant location and often incorporates different learning mediums to provide a range of educational opportunities (Moore, Dickson-Deane and Galyen, 2010). Taylor and Mohr (2001) created an online distance learning mathematics course which incorporated student-centred strategies such as informal language, relevant contextual materials and reflective practice techniques with the dual aim of improving student's mathematics knowledge as well as easing their mathematics anxiety. 90% of the users reported that the course had improved their mathematics confidence. Taylor and Mohr's research also proposed that distance learning may be more encouraging for students who are reluctant to discuss their past mathematics issues within a traditional classroom environment. The use of distance learning as a basis for the learning resource proposed in the objectives of this research (RO2.b) could therefore be promising as a method to assist learners with the Machine Learning domain. Online learning is reviewed in Section 2.9. Due to the prevalence of mathematics within this domain the two approaches of retesting and self-paced learning will also be considered for incorporation to build learner confidence and ease any anxiety relating to mathematics.

2.8.2 Confidence and Self-Efficacy

As well as anxiety with particular aspects of the subject domain, students may also have a lack of confidence in their ability to adequately learn the subject to a competent level. This strength and sense of self-belief in our learning capabilities is called self-efficacy (Bandura, 1977; Hattie, 2012). Table 2 outlines the differences between learners who possess a high level of self-efficacy compared to students who have low levels of self-efficacy.

As shown in Table 2, students who have a high-level of self-efficacy can have a different mindset to learning compared to learners who have lower self-efficacy, including a view of difficult material and troublesome tasks as a challenge rather than something out of their grasp and therefore avoidable. Students with lower self-efficacy are also more likely to suffer with low confidence in their skills, particularly after an event which they deemed a failure.

<p>Self-efficacy is the confidence or strength of belief that we have in ourselves that we can make our learning happen.</p>	
High self-efficacy	<ul style="list-style-type: none"> - Student sees hard tasks as challenges rather than tries to avoid them. - Student perceives failures as chances to learn and to make a greater effort or to look for new information next time.
Low self-efficacy	<ul style="list-style-type: none"> - Student is more likely to avoid difficult tasks, which are viewed as personal threats. - Student has low or weak commitment to goals. - Student sees failures as chances to dwell on personal deficiencies, obstacles encountered, or to deny personal agency. - Student is slow to recover a sense of confidence.

Table 2: The different characteristics of students with high and low self-efficacy (Source: Hattie, 2012)

Psychologist Albert Bandura proposed the concept of self-efficacy and hypothesized that expectations of our own personal efficacy determines whether and to what extent our coping behaviour will be initiated, how much effort we will expend and how long this effort will be sustained in the face of obstacles and adverse experiences (Bandura, 1977). Bandura advised that the level of confidence someone has in their own effectiveness can affect how they will attempt to cope in a variety of situations. Bandura's findings posit the importance of self-efficacy, particularly in the face of challenging circumstances, therefore it is important to build learners self-efficacy and their belief in their ability to achieve the learning outcomes. Bandura (1977) proposed a number of differing ways in which self-efficacy can be enhanced, including self-directed mastery opportunities which provide opportunities to develop coping mechanisms and behaviours which can lessen personal susceptibility to stress. Alongside self-directed mastery, the opportunity for successful experiences relating to independent performance can reinforce the individuals feeling of competence, in turn boosting their self-efficacy (Bandura, 1977; Hoban and Hoban, 2004). Enhancing student self-efficacy is key to RO2.c and RO2.e. Inclusion of strategies to assist students in building their self-efficacy within the learning resource (RO2.b) will be explored as this may also help learners when they encounter the threshold concepts.

Hattie and Donoghue (2018) identify the importance of recognising that students have confidence that they have a feasible chance at attaining the success criteria outlined, for

example in a module. It is also important that the students see the value in the lectures and can relate them to prior learning and subsequent desired skills but do not feel unduly anxious about the skills they are being required to master (Hattie and Donoghue, 2018). One method to ensure students are not over-anxious about the module requirements is to specify initially within the module what successful learning in the lessons will look like and what the success criteria are. This will help reduce anxiety, potentially increase their motivation and build both surface and deeper understanding (Hattie and Donoghue, 2018) and may consequently improve their self-efficacy.

Within the field of Computing, specifically Data Science, a concerning finding from the recent study on the UK Data Skills Gap (Department for Digital Media and Sport, 2021) found that many students (45%) surveyed did not feel “well-equipped to carry out future data roles when they enter or re-enter the workplace” (GOV.UK, 2021, p.33). This finding implies a potential issue with the current educational provision within this domain and a need to investigate levels of confidence and self-efficacy of students partaking in such courses. Therefore, evaluation of differing cognitive strategies is a key focus of this research (RQ3). Entwined within the concept of self-efficacy is that of metacognitive knowledge and regulation which combine with suitable directed motivation (Luckin, 2018). These concepts will be covered in the following sections.

2.8.3 Metacognition

“Metacognition refers to the ability to reflect upon, understand, and control one’s learning” (Schraw and Dennison, 1994). This is a critical skill which enables students to self-monitor and self-regulate to advance a skill set (Scharff *et al.*, 2017). Self-efficacy and confidence (as discussed above) are pertinent to metacognition as students are inherently required to express a level of confidence when solving a problem posed to them and students’ self-judgements, especially of their confidence level, implores the learner to monitor their metacognition (Stankov, Morony and Lee, 2014).

The building of metacognitive skills is centred around enabling learners to reflect and analyse their thoughts to be able to put their learning into practice (Drummond, 2009). Students also need to have some understanding of how their mind functions, such as how they perceive

their ability to perform important cognitive tasks including remembering and problem solving (Downing *et al.*, 2007), tasks pertinent to Computer Science education.

There are a number of methods based around the concept of improving student's metacognition and ability to self-regulate their learning (discussed in the following section), such as the use of problem-posing education. Problem-posing education was recommended by Freire (1970) in opposition to banking education, where education becomes an act of depositing information (Freire, 2004). Problem-based education bolsters students critical thinking and "stimulates true reflection and action upon reality" (Freire, 1970, p.84). Teaching metacognitive techniques within the context of a specific course discipline has shown impressive results in that learners who participated in the study achieved significantly higher grades than non-participants (Volet, 1991). The metacognitive strategies applied included social interactions which promote the transfer of higher-level thinking skills based upon advice and guidance from the instructor (Volet, 1991). Methods such as problem-posing education and strategies enabling the learner to reflect on their current understanding of the domain are potential methods for inclusion in the learning resource (RO2.b, RQ3).

2.8.4 Self-Regulated Learners

Self-regulation entails the "monitoring and managing of one's cognitive processes" as well as control over factors such as emotions, behaviour and environment pertaining to learning (Nilson, 2013a). Taking responsibility for your own learning is a foundational element of higher education (Drummond, 2009). However, not all adult students are automatically self-directing and may have difficulty learning to become so, as self-direction requires students to learn new skills, modify their learning habits and increase their self-confidence (Kegan, 2018). Although cognitive self-regulation can be difficult to achieve there is a growing body of evidence which indicates that this skill can be taught and it has been shown that students who apply self-regulatory skills achieve higher grades within the content domain in which these skills apply (Boekaerts, 1997).

There is a growing consensus on the importance metacognition plays within the role of self-regulation, as students who have more sophisticated levels of metacognition, specifically related to the subject matter domain, demonstrate greater strategy use and are often better

problem solvers (Boekaerts, 1997; Hattie and Donoghue, 2016). Therefore, it is important to develop student's metacognitive skills in order to help them become greater self-regulated learners (RO2.c).

Comprehension of any deficits students have in their self-regulating techniques may help to inform pedagogical methods to aid students, particularly in overcoming threshold concepts (RQ3). For example, helping students to reflect on their learning may highlight specific content areas which require further attention. Hattie and Donoghue (2016) suggest varying strategies related to self-regulation including elaboration and organisation, concept mapping, monitoring the strategies used, elaborative interrogation and metacognitive methods. The overall goal for students to become self-regulated is that "they know what to do when they do not know what to do" (Hattie and Donoghue, 2018, p.106). Inclusion of self-regulated approaches within the domain content will enable learners to develop these skills to equip them with the relevant strategies to persist when learning challenging content.

2.9 Online Learning

The proliferation of technology within learning has enabled various new learning techniques to be trialled. This technology has been shown to improve learner self-monitoring through the use of techniques such as quizzing with immediate feedback and gamification. These techniques have also been linked with motivation through new insights into the human reward system, research demonstrates that motivation provided by games of chance generate additional dopamine (Jones, 2010). Due to the pervasiveness of the technology and the potential benefits related to its use, such as alleviating mathematics anxiety (as discussed in Section 2.8.1), an examination of the literature pertaining to online learning has been completed to identify its potential to accommodate the learning resource outlined in the research objectives (RO2.b).

There are a number of different terms used for this type of learning including e-learning, distance learning and online learning. These terms are often used interchangeably however, these can be delineated. Distance learning can be used to describe the provision of education to those who are geographically distant, online learning is described as access to learning experiences through use of some form of technology and e-learning is more contentious in its

definition but can be characterised as applications, websites, objects and programs which provide a learning opportunity (Moore, Dickson-Deane and Galyen, 2010). This type of learning provision, when executed well, can improve student learning and achieve enhanced learning outcomes (Oliver, 2001). Therefore, this technique has potential to host the learning resource outlined in RO2.b (p.6), especially as the enhanced learning outcomes could relate to the identified barriers (RO1, p.6) and the aim to improve student's metacognition (RO2.c, p.6).

2.9.1 Massive Open Online Courses (MOOCs)

Massive Open Online Courses (MOOCs) have been one of the most notable trends in higher education over the past few years and relate to global, usually video-based content, problem sets and forums released through an online platform enabling a high volume of participants to partake (Baturay, 2015). The term first arose in 2008 and can be based on Stephen Downes and George Siemens connectivist distributed peer model (Baturay, 2015). Connectionism is a form of networked social learning that requires arrangement of a complex and diverse set of ideas into a network to form specific information sets, with the core skill being able to identify connections between information sources to facilitate continuous learning (Duke, Harper and Johnston, 2013).

There are a number of high profile MOOC platforms including Coursera (2021) which was founded in 2012 by Stanford University professors Andrew Ng and Daphne Koller (who are both highly respected scholars within the field of AI), Udacity (2021) and edX (2021) which hosts online courses created by institutions such as MIT and Harvard. These platforms host a number of courses in Computing, specifically playing a key role in educating the global workforce in AI skills (Perrault *et al.*, 2019). Courses offered include Introduction to Machine Learning offered by both Udacity and Coursera and courses such as Machine Learning with Python offered by edX. Average enrolment on MOOC modules is around 43,000 students, however only 6.5% usually complete the course (Jordan, 2014). This finding indicates a potential issue with retention in these courses which may potentially be linked to the pedagogy and situational environment linked to distance learning.

2.9.2 Online Pedagogy

There is some contention relating to best practice for appropriate pedagogy for online learning, there is currently a lack of cohesive, universal principles for this type of teaching. However, it has been recognised that the inherent variability of online education means that online instructional strategies may be more effective when tailored to specific educational contexts (Steele, Holbeck and Mandernach, 2019). However, there has been some generic suggestions relating to the instructional design approach. Oliver (2001) suggests a constructivist approach to instructional design for online learning and the inclusion of pedagogical goals inherent to this learning approach, including embedding learning within relevant contexts, the use of multiple modes of representation and motivating self-awareness in the knowledge construction process. Typical online courses are based around a weekly structure where students can access the relevant learning materials in their own time, materials and activities include quizzes, short videos and forums (Baturay, 2015).

Online learning tools are often incorporated into a blended or flipped classroom. This technique is an inversion of the usual pedagogical model typically employed with the 'traditional university lecture' (Forsey, Low and Glance, 2013). Blended learning entails a mixture of face to face contact time with the lecturer alongside online course content (Concannon, Flynn and Campbell, 2005). The aim of these teaching models is to shift the face to face engagement away from traditional lectures and instead to various seminars and symposia which actively engage the students in the process of discovery and consolidation of knowledge (Forsey, Low and Glance, 2013). There is existing evidence that in comparison to fully face to face modules or fully online modules, the blended learning approach is preferred by learners (Forsey, Low and Glance, 2013). This preference for this type of learning suggests that there may be a greater shift to this type of learning within the near future to meet the everchanging needs and requirements of students.

The innovation of higher education influenced by online courses is continuing and has been accelerated by the shift to online distance learning due to the COVID-19 pandemic (World Health Organization, 2021). Penprase (2018) suggests that the fourth industrial revolution and continuing integration of online learning will result in high quality, synchronous face to face learning alongside online technologies to enable learners to rapidly build skills and knowledge asynchronously. Therefore, consideration of the learning resource (RO2.b) being hosted

online will be reviewed upon analysis of the current AI provision to determine the suitability of this methodology.

Summary

The literature review has highlighted the pressing importance of investigation of the current education landscape pertaining to AI, specifically Machine Learning. As this technology becomes more ubiquitous there is an increased need for individuals skilled in this area and these individuals must be equipped with the relevant knowledge to become effective, responsible practitioners. This includes the embedding of ethics within AI curriculums and to continue to increase representation within this domain to promote and strengthen public trust in this technology. Reviewing the current educational provision will enable insight into the type of content taught on these modules and potentially any barriers or aspects which might affect student attainment (RQ2, RO2.e, p.6).

This review has pinpointed potential learning models which may be applicable to this domain including the Hattie and Donoghue model, of particular relevance to the objectives of this study is the 'will' input. This focus on equipping students with the mindset and strategies to overcome potential barriers is one of the key questions and objective of this research (RQ3). Assisting students in learning to learn alongside learning the content of this domain by improving their metacognition and helping them to become more self-regulated may assist them in persevering when they encounter challenges such as threshold concepts. Identification of the threshold concepts (RO2.a, p.6) and building the pedagogical content knowledge for Machine Learning (RQ1) will enable lecturers to plan for the difficult topics within their curriculum and implement strategies to ease students through these troublesome areas.

One of the potential barriers which arose through this analysis of existing research is the potential for difficulty with the mathematical aspects of these courses possibly from mathematics anxiety. To have a full understanding of this domain it is important that learners have some foundational understanding of mathematics and statistics. This baseline knowledge could be an issue due to the lack of mathematics literacy and the prevalence of mathematics anxiety. Therefore, it is important to discern the level of mathematics knowledge and anxiety within these student cohorts. Even if mathematics anxiety is not found to be of

great concern within student cohorts it is worthwhile to incorporate some form of mathematics and statistics provision within the proposed learning resource outlined in RO2.b to try to alleviate any gaps in student's educational background. As identified in the literature, potential strategies could include retesting, self-paced learning and potentially distance learning as these have also been shown to improve student self-efficacy and self-regulation. Improvement of this specific set of skills alongside metacognition could potentially equip the learners with a wider range of strategies to tackle the challenging aspects of these courses, such as the threshold concepts (RQ3).

The lack of pre-requisites on Computer Science courses and therefore AI courses, as these are usually situated within the Computing department may also potentially be a barrier for student attainment. Although the lack of prior knowledge requirements for admission onto such courses may be a barrier to learning it can also be a benefit to widen participation on such courses. Therefore, the creation of a learning resource to attempt to lessen these educational gaps may be a better solution than increased formalisation of pre-requisites. Identifying the learner's mental models of this domain early on within the module should highlight to lecturers any misconceptions or areas for expansion.

Potential recommended practices to trial within the AI education domain have been identified by the review of existing Computer Science education literature, including the importance of active and problem-based learning activities founded upon the constructivist learning theory. Allowing students to explore varying Machine Learning models before delving into syntax and programming details may help build the students mental model before considering the less high-level concepts of this domain. The minimisation of mathematics notation, particularly as a method to introduce students to specific Machine Learning models may also be a best practice consideration and instead replacing this with visualisations to introduce them to the high-level concepts. Producing a framework for an introductory Machine Learning course (RO2.d, p.6) will provide lecturers with guidance on teaching this domain. This was identified as something lecturers would like to aid with their curriculum design for this domain alongside additional resources, particularly online resources to assist them with their teaching.

In the following chapters an investigation into the current education provision relating to AI is presented, including an analysis of online information pertaining to the courses offered at HE

level and a number of case studies at a range of institutions. The specific areas of investigation encompass the issues raised within this literature review and in the fulfilment of the research objectives RO1 and RO2 as outlined in Chapter 1 (p.6).

Chapter 3. Methodology

3.1 Introduction

This chapter details the data sources used to explore the current offerings in AI within higher education, with the aim of identifying the barriers and difficulties which may impact student attainment (RQ2, RO1, p.6). The findings from this data collection will inform many aspects of the proposed learning resource (RO2.b, p.6), including the content, based upon the initial threshold concepts identified (RO2.a, p.6), as well as strategies to improve student metacognition and self-regulation (RQ3, RO2.c, p.6). A mixed methods research approach was chosen which consists of the logical integration of both qualitative and quantitative approaches (Loui and Borrego, 2019) aiming to provide a more complete and comprehensive understanding of AI education. This approach can provide multiple views of the subject under investigation, therefore “increasing the usefulness and credibility of the results found” (Cohen, Manion and Morrison, 2018, p.33). This form of methodological triangulation enables alignment of multiple perspectives (Salkind, 2010) which is pertinent to this research to determine both the lecturer and student view of education within the AI domain. Triangulation of the data collection methods will assist in providing a definitive answer to the research questions and assist in achieving the research objectives.

The methods of data collection are outlined in this chapter, including the case studies undertaken at three different universities, the manner in which the data was collected, and the steps taken to format and structure the data before analysis. The chapter concludes with a discussion on the qualitative and quantitative methods used to analyse the data.

3.2 Data Sources

Determining the potential data sources was informed by the overall research aim of gathering as much information surrounding the current offerings in AI education to ascertain how this subject is currently taught and to discover the barriers to learning and teaching this topic. This round of data collection will then form the basis for the next stage of this research, informing the creation of an AI learning resource (RO2.b). In order to gather as much information

pertaining to the current HE AI offerings, a range of data collection methods were employed including an online review of modules, questionnaires, interviews and observation.

Before commencing any of the data collection techniques, formal ethical approval was obtained from the university ethics board (application number: 18-ALL-039) to ensure that the proposed research met university regulations and that all aspects of concern had been considered and where possible mitigated against. Alongside formal ethical approval, each data collection method underwent due diligence to ensure that all potential ethical implications had been considered, these will be discussed in the respective method sections.

3.2.1 Online Review of Modules

To determine how ubiquitous AI is within the computing curriculum an online review of modules was completed. This data collected from this method contributes to the answering of RQ1 and attainment of RO2.a and RO2.d. The criteria for this review required the institution under investigation to be classified as higher education and had to offer some form of course which included the keywords: “AI”, “Machine Learning” and “Deep Learning”, courses which were classified as “Data Science” were also reviewed. Both the overall degree programme (e.g., MSc Artificial Intelligence) and the specific modules (e.g., Introduction to Machine Learning) were reviewed. The data collected from the institutions which offered some form of AI course included:

- The level the programme was offered (e.g., undergraduate / postgraduate)
- Whether the programme was optional or compulsory
- Module pre-requisites
- Module content
- Module structure
- Learning outcomes
- Assessment

Where available, this information was collected from the university website. In most cases this information was publicly and freely available. A number of the institutions also listed the personnel involved in the delivery of the specific modules, which helped inform the contact list for the questionnaires and interviews.

To provide a systematic approach to this method of data collection, the review focused on UK universities, as well as the US and wider Europe which were present on top university lists and provided some form of AI course. Constraining the geographical location of the universities under review and systematically evaluating the universities present on league tables enabled a bounded list of universities to analyse and determination of what the top-rated universities are offering in terms of AI education provision.

There were a number of aims for this form of data collection including the identification of core concepts which are deemed important to teach on an introductory course to assist in fulfilment of RO2.d (p.6). Determining the differing learning methodologies used when teaching this domain and the type of educational background students are required to have before commencing a course in AI were also aims of this form of data collection.

There were some limitations to the data collection, including the narrow scope of the search criteria. The interdisciplinary nature of AI means that this topic is often taught on non-computing courses, for example Business degrees may offer some form of instruction on AI. However, computing departments and the courses they offer were the main focus of this method. Searching top university lists may also have restricted the variation of universities prevalent within the findings, however this provided a starting point. As well as structure to this form of data collection.

Upon collection of the data and before analysis could begin, the findings were categorised by the type of university, for example UK universities were labelled as either part of the Russell Group, a collection of research intensive universities (Russell Group, 2021) or were part of the ranking of top universities. The findings were also put into categories of undergraduate or postgraduate, module pre-requisites and content.

3.2.2 Questionnaires for Lecturers

The use of a questionnaire for lecturers was to determine what they perceived as the difficulties faced by students when learning this domain, as well as the difficulties lecturers face when teaching AI (RQ2, RO1). This data collection method was also used to further identify how AI is currently being taught, with content relating to taught material and pre-requisites (RQ1). Potential participants for the questionnaire were identified when reviewing

online information relating to AI modules, through their inclusion on the webpage of an AI module. Upon identification as a potential participant, an email was sent which detailed the aim of the research and included the link to the online questionnaire (Appendix A). The potential participants were informed that the questionnaire was anonymous. The questionnaire was anonymous because of the ethical consideration and the design. There was also an absence of questions which asked potential identifying details such as institution name. The respondent was also informed that by completing the questionnaire they were consenting to their data being used as part of this study and were provided with an email address in which to ask further questions. Alongside the questionnaire, the identified lecturers were also asked about potentially participating in an interview to further discuss their experiences teaching AI.

The questionnaire content is based around six general themes: (1) the general *demographics* of the students who participate in the module the lecturer teaches. The questions based upon demographics included whether the students are undergraduate or postgraduate, whether the course was compulsory and a general description of the gender demographics. (2) Understanding the course *pre-requisites and student educational background* was important to comprehend the expectations and requirements for students studying these types of modules and to ascertain some of the issues students may be facing. This data may also illuminate some of the broader issues within widening participation and diversity within these types of courses by highlighting any gender imbalances or predilection to individuals with a particular educational background. (3) Ascertaining the *content* being taught on courses within the AI domain will enable cognition of the aspects of AI being taught, for example is there more of a push on Deep Learning than more focused applications such as robotics? (4) The questionnaire contained questions relating to what the lecturers felt was *core knowledge*, for example they were asked to note topics which they think are essential to teach. It was also pertinent to ask, in their experience which aspects of the module the students struggle with. Theme (5) related to *pedagogy* and the strategies the lecturer employs to aid their teaching and (6) *feedback* they receive from students on their module.

Determining assessment methods and additional resources lecturers use alongside teaching strategies gave insight into the current approaches to teaching this topic. Acquiring information relating to the feedback the lecturers receive from students who have

participated in their module enabled a general overview of how students were finding such courses and an insight into student satisfaction on such modules (R02.e, p.6). However, if the lecturers received negative feedback from students, they may have been wary about disclosing this within the questionnaire.

It was important to carefully consider the content and the type of questions to be included within the questionnaire as there is a risk, especially when categorising responses in “reducing the participants to data objects rather than agentic people” (Cohen, Manion and Morrison, 2018, p.120). There was a balance to be struck between ensuring the participants fully understood the information being asked of them and ensuring that any prompts or instruction did not unduly influence their potential responses. Therefore, any information to assist the participants in completing the questionnaire was only related to how to register their response.

3.2.3 Interviews with Lecturers

Alongside the questionnaire, participants were asked about a potential interview. The correspondence contained details of the research and why they were being approached as potential participants and what their presence in the study would entail. Like all of the other methods conducted for this research, participation was optional and potential interviewees were offered the option of a discussion session prior to deciding whether they would like to participate, if they felt they required further information. Participants were also advised that upon completion of their interview the results would be anonymised, and any findings reported would not be traceable to them.

The interviews were conducted in a semi-structured, open-ended style to gain the interviewees’ perspective via *how* and *what* questions (Brenner, 2006). Questions referred to the student cohort, module delivery methods, course pre-requisites, module content and topics which they felt students struggled with (questions in Appendix B). The questions map to the overall aim of the research in determining a best practice framework as well as RQ1, RQ2 and RQ3. All of the interviews were either conducted in person face-to-face or via video call. This organisation allowed for the observation of body language and the creation of a safe environment where participants could speak freely (Yeo *et al.*, 2014).

The interviews were also used to identify the lecturers' perception of key threshold concepts and troublesome knowledge within the subject domain of AI (RQ2, RO2.a). This identification is usually the task of the subject instructor and is based on their experience of teaching and working in the particular domain of interest (Davies, 2006; Serbanescu, 2017). Threshold concepts can also be identified through course participant enquiry, for example, through discipline specific interviews (Serbanescu, 2017). Interviews with course participants were out of the scope of this research. Instead, participant views on threshold concepts were gathered through questionnaires and the one-minute paper as part of the case studies discussed below.

3.3 Case Studies

During the interview process, subjects were broached with the idea of furthering participation in the research by contributing to a case study. There was some interest from a number of the lecturers, however this was only the first stage in obtaining access to the institutions. In many cases gatekeepers were present which required their permission in order to further the research. When approaching the gatekeepers, it was important to clarify the scope of the project, outlining to the relevant parties the aims, methods, duration and data procedures so that they had a full picture of the research being undertaken (Cohen, Manion and Morrison, 2018, p.134). Upon consultation with some of the gatekeepers from the interested universities, unfortunately some of the interested parties were unable to take part in the case studies as course administrators did not permit participation. However, three universities were able to participate. One of these was within the Machine Learning module at Newcastle University, a Russell Group university (referred to as University A), and the other two were within universities ranked in the top 100 within the UK (The Guardian, 2020), these will be referred to as University B and University C going forward. The mixed type of these universities, alongside the data collected from the other study methods including the online review of modules (Section 4.3) will enable generalisability of the findings.

The overall aim of the case studies was to determine how AI was taught at that specific institution and to determine both lecturer and student opinions and experiences (RQ1 and RQ2). The main methods of data collection included questionnaires and the one-minute paper (Stead, 2005). Observation of AI lectures was also conducted to collect contextual evidence relating to the student and teacher experience in the classroom. The initial aim of the case studies was to complete all data collection methods at all three of the participating

universities, however, this was limited by the voluntary nature of the study and the level of data collection permitted at each institution.

Within the postgraduate Machine Learning module at Newcastle University (university A), pre- and post-module questionnaires were completed alongside observation and the one-minute paper. Data collection differed at the other two universities, one of the modules which was used as part of the case study was an undergraduate module in Artificial Intelligence (University B). The data collection methods employed within this module included observation and a post-module questionnaire. Observation was completed at the third university (University C) as part of the undergraduate Machine Learning and Computer Vision module. All of the lecturers from the case studies participated in the semi-structured interviews. The inability to complete all data collection methods at each of the universities participating was a drawback within this data collection method, however the range of methods used across this study and triangulation of the data mitigate against this.

The data collection methods conducted with students as part of the case studies were as follows:

3.3.1 Pre-Module Questionnaires

Before approaching students to participate in the study, time was taken within the first lecture to discuss the research with the students. They were informed of the aim of the study and that participation in the questionnaire was optional and anonymous. The pre-module questionnaire was designed to determine the students' educational background, including previous qualifications, mathematics and programming competency and their overall expectation of achievement level. Getting students to self-report on their mathematics knowledge and confidence within this domain is one of the key methods in identifying mathematics anxiety (Ramirez, Shaw and Maloney, 2018).

When designing the questionnaire there were three main areas of the research to be covered, these included background information relating to student demographics and educational background, student's prior knowledge of Machine Learning and teaching and learning strategies. The questionnaire (Appendix C) starts with questions pertaining to demographics including categories for age and gender to determine who is undertaking courses within this

domain. As the questionnaire was to be employed within a postgraduate level module, determining the students educational background was important to discern whether they had previous experience within computing or mathematics. This set of questions were pertinent to RO1 (p.6) to identify whether a lack of prior experience in either mathematics or computing may be a potential barrier (RQ2). The participants were also asked to rate their confidence in both their mathematics and programming skills, these questions are pertinent to RO2.c (p.6) to determine the need for improvement of student's metacognition and related to RQ3 (p.6). It has been shown that self-efficacy and self-regulation are intrinsically linked to mathematics anxiety (Jain and Dowson, 2009). Efficacy plays a large role in expectations and coping strategies when faced with potential obstacles and adversity, the greater the recognised individual self-efficacy, the more resolute the student will be in actively attempting to overcome difficulties (Bandura, 1977).

Questions relating to Machine Learning were included to determine whether students undertaking courses within this area had a pre-existing mental model of this domain or any misconceptions. These questions relate to RO1 (p.6), to determine whether a learner's mental model/accuracy of mental model can be a potential barrier to learning. Students were also asked how interested they were in learning this topic, this was an attempt to gauge student motivation levels as well as how confident they were in their ability to do well. The inclusion of these questions was based on Bandura's method of self-motivation where self-motivation provides performance evaluation standards and where "perceived negative discrepancies between performance and standards create dissatisfactions that motivate corrective changes in behavior" (Bandura, 1977).

To conclude the Machine Learning section, the students were asked to reflect and note down whether they expected to find the theoretical aspect of the module harder or the practical exercises. The inclusion of this question was to determine whether the student view correlated with that of the lecturer relating to which aspect of the domain proves most difficult to learn.

The purpose of the teaching section was to determine which skills and resources students identified as the most useful in helping them learn in this domain. The participants were asked which skills they had from previous learning they feel would aid them within the Machine

Learning module. They were also asked to rate a list of resources in order of importance to their learning. These resources included lectures, written material and practical exercises. The inclusion of these resources was to determine out of the teaching strategies employed by the lecturer which they rated the most important (RQ1). As well as determining useful teaching strategies, the students were also asked to rate, in order of priority, where they would most likely go for additional support. Options for this question included the module leader, demonstrators for the practical sessions, peers and resources such as textbooks and online information. This question was included to determine which resources the students prioritise and where they are looking for additional support. As with the other questions which included outlined categories, an 'other' section was included so that the participants could enter in their own response if this was not covered within the pre-defined options.

Finally, the students were asked to tick all of the learning strategies they proposed to use within the Machine Learning module, as shown in Figure 7. Alongside practical skills such as note taking and the use of online guidance were more holistic strategies such as reflection, goal setting and self-evaluation to determine whether the students were using techniques which would improve their self-regulation.

d) In order of priority, rate where you would most likely go for additional support: <i>Please rate these from 1-6, with 1 being the first choice</i>	
Support Option	Rating (1-6)
Module leader	
Demonstrator	
Peers	
Textbook	
Online	
Other (please state):	

e) Which of these learning strategies do you propose to use in this module: <i>Please tick all which apply</i>			
Note taking <input type="checkbox"/>	Study group <input type="checkbox"/>	Practical exercises <input type="checkbox"/>	Quizzes <input type="checkbox"/>
Textbooks <input type="checkbox"/>	online guidance <input type="checkbox"/>	Critical thinking <input type="checkbox"/>	Reflection <input type="checkbox"/>
Goal setting <input type="checkbox"/>	Planning <input type="checkbox"/>	Self-evaluation <input type="checkbox"/>	
Other (Please state):			

Figure 7: Questions 3d and 3e from pre-module questionnaire

After introducing the purpose of the research and explaining my presence within the module, the questionnaire was handed out alongside the information sheet and research consent form in paper format to the students at the end of the lecture. The students were also informed that I would be present within the following practical session if the students required any further information or had any specific questions or concerns to be addressed.

3.3.2 Observation of Lectures

Observation was chosen as a research method as it has a number of advantages, including the “potential to yield more valid or authentic data than would otherwise be the case with mediated or inferential methods” (Cohen, Manion and Morrison, 2018, p.542). Observation can also provide data on “interactions, processes and behaviours that goes beyond the understanding conveyed in verbal accounts” (Ritchie *et al.*, 2014).

Observation was carried out at all three of the institutions participating in the case studies, as this data collection method was being used across institutions a standardised observation guide was prepared for use within the sessions (Appendix D). This ensured that specific criteria were assessed within all observed sessions and enabled the results to be comparable. The observation guide was comprehensive and contained numerous categories such as the purpose of the lecture and the topics to be covered, the physical environment and attendance. The main focus of the guide was to log the sequence of events and activities that occur during the lecture as well as interactions between the lecturer and students and between peers. This data will be relevant in determining current best practice within this domain (RQ1). The guide concluded with space to reflect on the observed session and note any perceptions or further lines of enquiry. Within all the of the universities the role of observer was as “observer-as-participant” (Cohen, Manion and Morrison, 2018, p.543). Group activities within the lecture were not participated in and observation was as unobtrusive as possible. However, the role of the observer as a researcher is clear and overt to the student cohort.

3.3.3 One Minute Paper Technique

Within the postgraduate Machine Learning module at Newcastle University (University A), the one-minute paper method was employed to gauge how the students were finding the module, particularly to identify any difficulties they were having with certain topics (RQ2) and to

provide insight into potential threshold concepts (RO2.a, p.6). This data collection method was only employed at university A as there was greater flexibility in researcher participation within this module and greater scope for use of data collection methods due to the cooperation of the university.

The one-minute paper is an uncomplicated, adaptable technique to gather student feedback relating to their education experience, it can be used anonymously and is based on simple questions based on their learning experience (Chizmar and Ostrosky, 1998). The one-minute paper has many benefits as a data collection method, including the potential for students to ask specific questions which they may have been reluctant to ask within the lecture, it also promotes student reflection (Stead, 2005).

The students were informed that completion of the one-minute paper was optional and that any responses were anonymous. It was communicated to the participants that the responses from the one-minute paper would be used to inform the revision session and that any specific questions asked would be answered within the next lecture. As the module was block taught, the one-minute paper was handed out for completion at the end of the two taught weeks which covered the main assessed content (Appendix E).

The aim of the use of the one-minute paper within this study was to gain an understanding of specific topics within the module that the students were having difficulty learning in an unobtrusive and least time-consuming manner. As shown in Figure 8 the paper was designed to be completed in less than five minutes and was adapted from the more commonly used one-minute paper templates to match the requirements of this study. The one-minute paper used within the Machine Learning module contains four questions, the first question asks the students to reflect on the topics covered in the lecture series of that week and determine which they found the most difficult. As the module was block taught, the lectures were fairly intense and covered a variety of content, to help the students reflect on the topics they had learnt, a topic list was included. For the second question, the participants were asked to note any subject which they were still unsure on. These two questions were included to discern which topics were proving the most troublesome for the students and to be able to feed this back to the module leader so they could provide further resources or instruction.

Week 1																													
Topics Covered this week:																													
<ul style="list-style-type: none"> • Introduction to Machine Learning (data representation, supervised /unsupervised learning, overfitting) • Maths Primer for Machine Learning • Linear Regression • Probability Theory • Gaussian Distribution/Gaussian Mixture Models • Maximum Likelihood Estimation • Logistic Regression Classifier • Deep Learning • Softmax Logistic Regression • Multilayer Perceptron • Autoencoder • Dropout • Stochastic Gradient Descent • Convolutional Neural Networks • Recurrent Neural Network 																													
Which topic(s) did you find the most difficult this week?																													
Which topic are you still unsure on?																													
Any specific questions?																													
How confident do you feel in what you have learnt this week? <i>Please circle a number</i>																													
<table border="1" style="width: 100%; text-align: center;"> <tr> <td>1</td> <td>2</td> <td>3</td> <td>4</td> <td>5</td> <td>6</td> <td>7</td> <td>8</td> <td>9</td> <td>10</td> </tr> <tr> <td colspan="5">Not confident</td> <td colspan="5">Very confident</td> </tr> </table>										1	2	3	4	5	6	7	8	9	10	Not confident					Very confident				
1	2	3	4	5	6	7	8	9	10																				
Not confident					Very confident																								

Figure 8: The One-Minute Paper (Week 1)

The third part of the one-minute paper gave the participants the opportunity to ask the lecturer any outstanding questions they had. The questions the students asked would enable cognition and identification of gaps in the students' knowledge as well as highlight specific content knowledge which the students had questions about. Finally, the students were asked to rate their confidence in what they had learnt that week out of a scale from one to ten, with ten being exceptionally confident. Asking the students to reflect on their confidence was included as confidence can be an important predictor of performance (Stankov, Morony and Lee, 2014) as well as inform on student cognitive bearing, self-beliefs and motivation (Stankov *et al.*, 2012).

3.3.4 Post-Module Questionnaires

This questionnaire was designed for completion at the end of the module, with two of the institutions participating in this data collection method (questionnaires in Appendix F). The post-module questionnaire was offered to participants on the module at Newcastle University

(University A) and within the undergraduate Artificial Intelligence module at University B. The questionnaire was hosted online for ease of use for the lecturers so that they could include the link to the form within their lecture slides. The students were informed that completion of the questionnaire was optional, and all data was anonymous.

The aim of this data collection method was to determine how the students found the experience of completing a module within this domain and potential areas for improvement (RQ1, RO2.e, p.6). The majority of the main content for the questionnaires were identical for the main points of questioning, however there were a few differences to ensure the questionnaires were tailored appropriately for the specific institution. As a pre-module questionnaire was offered at University A, lines of enquiry related to mathematics and programming confidence were not repeated, however these were included within the questionnaire for University B to determine educational background.

One of the objectives of the questionnaire was to determine which topics the students found the most challenging (RQ2). Students were presented with a list of the main topics covered within their module and asked to rate how challenging they found them on a scale from one to five (five = most challenging). They were also asked to select which topics they are still unsure on and whether they found the theoretical or practical aspects more difficult. These questions would provide insight into potential threshold concepts (RO2.a, p.6).

The next thematic section of the questionnaire related to study strategies and resources for learning which the students used during the module. Determining which strategies for learning the students are using is important to understand how they are approaching their learning within this domain and to understand the type of resources they are using to aid them and provide further instruction. Preference for a particular type of learning resource will assist in choosing the technology for the learning tool to be created (RO2.b, p.6).

Participants were also asked to rate on a scale of one to five (five = most challenging), how confident they felt about applying Machine Learning/Artificial Intelligence techniques upon coming to the end of their course. This question was asked to provide some indication of how confident they felt in applying the new skills they have learnt. Students were also provided with the chance to leave any additional comments about their experiences in case they had a

concern or experience to share which was not covered within the questionnaire. Finally, the students were asked two questions based on their demographics, these included their age and gender. As these are protected characteristics, the participant had the option, as with all of the questions on this form to leave this option blank or select 'Prefer not to say.'

3.4 Analysis Methods

3.4.1 Statistical Methods

A number of different statistical techniques were used for the data collection methods to analyse the data and to test for reliability. Before determining the statistical tests to be undertaken, it was important to ascertain the scale of measurement, e.g., nominal, ordinal, interval or ratio of the data as well as the data distribution as this influences the types of tests which are applicable. The majority of the data can be categorised as either nominal or ordinal (Cohen, Manion and Morrison, 2018), therefore nonparametric statistical tests are appropriate. However, parametric tests are applicable for some of the data collected from the questionnaires. Both descriptive and inferential statistics were applied to our data, descriptive statistics have various applications including description of the characteristics of the dataset and testing to see that the variables do not violate any underlying assumptions before undertaking further statistical techniques (Pallant, 2016). Inferential statistics "strive to make inferences and predictions based on the data gathered" (Cohen, Manion and Morrison, 2018, p.727). Table 3 displays the proposed statistical methods for each data collection instrument.

Data Collection Instrument	Proposed Analysis Techniques
Online Review of Modules	Frequency analysis, thematic analysis
Questionnaires for Lecturers	Frequency analysis
Interviews with Lecturers	Thematic analysis, frequency analysis, inductive process
Pre-Module Questionnaires	Skewness and kurtosis, measures of central tendency, frequency analysis, crosstabulations, correlation, ANOVA, Chi-Squared Test for independence
Observation	Thematic analysis, frequency analysis
One-Minute Paper	Thematic analysis, frequency analysis, measure of central tendency
Post-Module Questionnaires	Measures of central tendency, frequency analysis, skewness and kurtosis, crosstabulations, ANOVA

Table 3: Proposed analysis techniques

The following sections describe in detail the specific statistical techniques to be applied from Table 3.

Descriptive Statistics

Determining the type of descriptive statistic techniques to apply required categorisation of the type of variables to determine whether our data was categorial or continuous as many of the methods of data analysis are not appropriate for either type of data. Once the variable type was determined, the appropriate statistical tests could be applied. The descriptive statistic tests applied included:

- *Skewness and Kurtosis*

Skewness and kurtosis are related to the distribution of the data and are usually measured for continuous data. These measures were applied on the post module questionnaire for University A (Section 3.3.4) on the scales relating to confidence and attainment. Determining the distribution of the data is an important step within the analysis process as this influences the type of statistical procedures which can be applied. A normal distribution contains the higher frequency of scores within the middle and the smaller frequencies towards either end of the extremes and is represented as a bell-shaped curve (Pallant, 2016). If the data is normally distributed then parametric tests can be applied, such as t-tests and analysis of variance, however, non-parametric tests will be required to perform analysis if the distribution is not 'normal.' The skewness and kurtosis provide information on the distribution, skewness relates to the symmetry as a normal distribution is symmetrical. Kurtosis is focused on the 'peak' of the distribution curve, measuring the steepness of the slope and the spread of the data surrounding the peak (Cohen, Manion and Morrison, 2018). A perfect example of a normal distribution would be a value of zero for both skewness and kurtosis, however, this is an exceptionally rare event (Pallant, 2016). Alongside analysing the skewness and kurtosis when assessing normality it is also recommended to investigate the shape of the distribution using a histogram (Tabachnick and Fidell, 2013).

- *Measures of Central Tendency*

Measures of central tendency for the purposes of this data analysis included mean, mode and median and are often referred to as summary statistics. The mean is broadly

understood as the most useful measure of central tendency, however it can be distorted by outliers and skewed data (Manikandan, 2011) therefore, it is advised to present the median as an alternative (Pallant, 2016).

- *Frequency Analysis*

For the data collected from the questionnaires, one-minute paper and the online review of modules, frequency analysis was employed using SPSS (IBM Corp, 2019). This method is employed to outline the data to enable identification of trends and important features in the data set (Qarabash, 2018). The output from the frequency analysis, like many of the descriptive statistics techniques is some form of visual data presentation (Cohen, Manion and Morrison, 2018). Frequency analysis will be a valuable tool in determining identification of trends, for example topics which were commonly identified as difficult (RO2.a, p.6) and confidence in particular concept areas (RO2.c, p.6).

- *Crosstabulations*

Crosstabulations are used to examine the relationship between two categorical variables, where, for example, nominal data such as male or female is contained in the rows and the ordinal data such as a five-point scale is in the columns (Cohen, Manion and Morrison, 2018). The crosstabulation is output and presented to display one variable in relation to the other. Crosstabulations are useful for identifying relationships which may not be immediately apparent and were used for some of the responses from the student questionnaires such as participant gender and their level of programming skill. Using crosstabulations on gender and self-identified programming skill level will identify if students of differing genders had a predisposition to self-identify at a particular skill level. This may indicate a particular barrier for members of the cohort (RO1, p.6) or an issue relating to confidence and metacognition (RO2.c, p.6).

- *Correlation*

According to Cohen et.al (2018, p.767), correlation is generally used to determine three questions about two variables or two sets of data:

1. ‘Is there a relationship between the two variables (or sets of data)?’ If the answer is yes then, then the following two questions are addressed:
2. ‘What is the direction of the relationship?’

3. ‘What is the magnitude of the association?’

(Cohen, Manion and Morrison, 2018)

The two most popular kinds of correlation statistics are the Pearson product moment correlation for use with interval and ratio data, and the Spearman rank order correlation for ordinal data. The correlation coefficient, for example the Pearson correlation coefficients (r), represents the direction of the relationship in a range of values from -1 to +1. A positive correlation, indicated by the ‘+’ sign, shows that as one variable increases so does the other one, a negative correlation (‘-’), indicates that as one variable increases the other decreases (Pallant, 2016). The size of the value, irrespective of the sign, indicates the strength of the relationship. Correlation will be applied to data points from the student questionnaires such as programming skill level and the importance students placed in the practical sessions. This will help identify the importance of educational background and the perception of useful learning strategies.

Inferential Statistics

Inferential statistics are different to descriptive statistics in that they “strive to make inferences and predictions based on the data gathered” (Cohen, Manion and Morrison, 2018, p.727). Although descriptive statistics reveal important information about data, inferential statistics are often considered the more powerful. The inferential statistic methods applied included:

- *Analysis of Variance (ANOVA)*

Analysis of variance can be understood as an extension of the t-test and can be employed to determine the differences between three or more groups. For example, one-way between groups ANOVA was used to determine if there was a difference in self-reported mathematics confidence for the different groups of mathematics attainment from the pre-module questionnaire for students. This will assist in identifying if educational background or lack of, particularly within mathematics can potentially be a barrier when learning AI (RO1, p.6).

Using SPSS to calculate the analysis of variance we are presented with two important outputs, the F-ratio which “is the between-group mean square (variance) divided by the within-group mean square (variance)” (Cohen, Manion and Morrison, 2018, p.782). Also

presented in the output is the p-value indicating whether there is a statistically significant difference between the means. However, the p-value only indicates whether there is a difference between the groups, further tests are required to confirm and identify what those differences are. The most commonly used post hoc tests include the Tukey test and the Scheffe test (Marshall, 2021). Effect size should also be considered, this value allows the determination of how much the independent variable has affected the dependent variable (Eddy, 2010).

- *Chi-Squared Test for Independence*

The chi-square test is used to explore whether there is an association between categorical variables and is a nonparametric test. This test will be applied on the student questionnaire data to determine any association between for example, gender and identified programming skill level. These results may have implications relating to widening participation within the AI domain and the potential barriers differing cohorts face (RO1, p.6). The test works by comparing the observed frequencies of cases which are present in each of the categories against the values expected if there were no association between the variables being measured (Pallant, 2016). Like other statistical tests it is important to check the assumptions for that test before applying it, for the chi-squared test, there is a requirement for a fairly large sample size with an expected frequency of one for each cell and an overall frequency of 80% for the majority of the cells.

Depending on the type of table, for example a 2 by 2 table (meaning that each variable has two categories), determines the output value we are interested in. The main value for inspection is the Pearson Chi-Square, however for a 2 by 2 table, Yates' Correction for Continuity should be used as it “compensates for the overestimate of the chi-square value” (Pallant, 2016, p.221). To be statistically significant, the value needs to be .05 or less.

3.4.2 Qualitative Methods

This section describes the qualitative data analysis methods employed to determine both lecturer and student opinions of the current offerings in AI education:

- **Inductive Process**

The majority of the qualitative data analysis, including analysis of the interviews, questionnaires with lecturers, observation and the one-minute paper, followed an inductive process. This is a bottom-up technique which progresses from the raw data, to explanations to theory (Cohen, Manion and Morrison, 2018). One of the purposes of the inductive analysis approach is to “develop a model or theory about the underlying structure of experiences or processes” (Thomas, 2006). This approach to some of the qualitative data analysis enabled a systematic approach to the development of themes and explanations of some of the phenomena present within our data.

- **Thematic Analysis**

Thematic Analysis, a “foundational method for qualitative analysis,” (Braun and Clarke, 2006) was used to analyse and identify patterns within the interview and observation data. The thematic analysis process involved coding and content analysis. Open coding was used to label the data which was then categorised using responsive categorisation which is an intuitive method in which categories are developed from the material instead of pre-determined (Cohen, Manion and Morrison, 2018). From the categories, additional statistical tests could be applied such as frequency analysis.

To address the potential for bias within the qualitative data analysis, inter-rater reliability analysis has been conducted for the coding of the lecturer interviews to determine the consistency of measurement (Fink, 2010).

3.5 Limitations of the Data

Often concerns are raised relating to the methodological approaches to pedagogic research, particularly within the HE sector as researchers within this area may be aligned to particular ideological values which may influence the methodological orientation (Stierer and Antoniou, 2004). To attempt to mitigate against any influence relating to the methodology of this study, a wide variety of data collection methods were used, this not only provided a wider pool of data for analysis, but also enabled greater perspectives from lecturer to student and even institutional views obtained through the online review of modules.

Tracking individuals, particularly from university A who participated in both pre and post module data collection may have provided further insight and contextualisation into their experience and resultant achievement within their educational journey learning AI. However, limitations governed by the participating institutions prevented this.

The UK-centric focus of this data collection, alongside the similar nature of the institutions participating in the case studies could be deemed too specific and inherent to the participating cohorts. However, this type of information is incredibly valuable to lecturers teaching this specific topic and as a consequence of the current lack of guidance related to pedagogy for teaching this domain, will provide guidance and help inform practice which is the main aim of this research. Evaluating and analysing all of the data collected through the various methods in a structured manner, using statistical tests was a way to make the evaluation more reliable.

3.6 Summary

In this chapter I have discussed the methods used for data collection to determine the current provision of AI education, alongside lecturer and student views of their experiences teaching and learning this domain. Methodological triangulation will enable analysis of multiple perspectives to address the research questions and objectives, for example data collected from lecturers, students and from the online review of modules can be combined to address the identification of the threshold concepts (RO2.a). This chapter has also outlined the methods for data analysis, including the different statistical tests applied and the rationale behind the application of these methods. The advantages and limitations of these data collection methods will be discussed further in Chapter 4 (p.81) and Chapter 7 (p.210) in relation to the results and findings from these methods. The next chapter presents and discusses the results.

Chapter 4. Current AI Education Provision and Experiences Within Higher Education

4.1 Introduction

This chapter outlines the results from the data collection techniques, statistical tests and qualitative analysis as discussed in Chapter 3: Methodology (p.61). The results from each of the data collection methods are presented alongside a discussion of the key themes and findings relevant to the study, i.e., identification of the barriers which might impact student attainment in Machine Learning courses (RO1, p.6). The key themes which have been determined from the data analysis will then be broken down and discussed individually alongside the implications for this research at the end of this chapter.

As discussed in Chapter 2, Section 2.3 consideration of the analysis of data can present a variety of ethical challenges. Potential bias within the data collection and the analysis and communication of results influenced many aspects of the work discussed within this chapter. Consideration was given to how to clean the data, how representative the data is of the population under investigation, with benchmark statistics provided when discussing protected characteristics. Figures are also provided relating to the number of respondents for the data collection methods to ensure transparency.

4.2 Study Context

The focus of this research was to inform on the issue's lecturers are facing teaching some form of AI, as well as the barriers to learning for students enrolled on such a course (RO1, p.6). Particular issues of interest were the presence of mathematics anxiety or low confidence in technical skills and the impact this may be having on student confidence and self-efficacy (RO2.c, p.6). One of the main aims of this particular round of research was to gain an overview of the current education offerings within higher education in relation to AI. For example, at what level is AI offered as a subject, undergraduate or postgraduate? What constitutes an introductory course and which topics are covered within this? (RO2.d, p.6) One objective of this research study was to create a learning resource (RO2.b, p.6) to be used by both lecturers and students. Identifying which resources are currently being employed is important to

determine which strategies both lecturers and students are currently using to gain insight into an appropriate methodology for the learning resource to be created.

4.3 Online Review of Modules

The online review of modules encompassed a systematic search of both the top 10 worldwide institutions for Computer Science (Times Higher Education, 2021) as well as the top 10 UK universities (Complete University Guide, 2021). This review was first undertaken in 2018 to help inform other data collection methods, however this method was updated in 2021 to determine if there had been any change in modules offered throughout this period. Three of the universities analysed were present in both of the top 10 lists. The heterogeneity of institutions present within the ranked list required additional universities to be analysed to build a broader picture of the current provision of education. Overall, 25 universities were subjected to an in-depth review of their current offering relating to AI.

The results reported in this section are from the publicly accessible information available through the university websites, therefore the institutions are named. The information collected was available within the period of 2018-2021 and was to the best of my knowledge an exhaustive list of modules offered and the information gathered was as stated at source at the time of data collection.

Over the course of the research period, a number of institutions increased their education provision relating to AI. The updated review of universities from the top 10 lists (Complete University Guide, 2021; Times Higher Education, 2021) identified that 16% of the 25 universities increased their offering of AI modules, providing not only introductory courses but also standalone modules in Deep Learning. Alongside the increase in modules available within this domain, many of the universities now offer degree specialisations in Artificial Intelligence.

The main focus of this analysis was to determine how Machine Learning is currently being taught. Modules which fall under the umbrella of Artificial Intelligence were included as Machine Learning is a sub-domain of this discipline and Machine Learning is often included within the content. Deep Learning can be categorised as a sub-discipline of Machine Learning; therefore, these modules were also subjected to analysis (refer to Figure 1, p.1 for an overview

of the domain). Including Data Science modules was also important to build a picture of the current educational landscape as Machine Learning techniques such as logistic regression (Cramer, 2002) are often taught as a means of analysing data. The results from this collection method are presented below, the findings from each of the module categories are presented first, followed by a summary of all of the main findings.

4.3.1 Universities Under Review

Determining the current provision of education in relation to Machine Learning required the inclusion of a range of differing universities. As part of the systematic approach to this form of data collection, international institutions were included. These institutions were determined through their inclusion in the top 10 universities in the world for Computer Science (Times Higher Education, 2021). However, their presence on this list also indicated a certain level of prestige as 70% of the universities in the list were international and included well renowned institutions such as Stanford University (2022). This indicates that the UK may be lagging behind in terms of Computer Science education compared to international competitors. There is scope in this method for further analysis of international provision to give a broader picture of how they teach their Computer Science courses to determine how they became such prestigious universities for this subject. The international universities analysed did give a good indication of the prevalence of AI within their Computing departments and how they approached the teaching of this subject. All of these institutions offered some form of AI module or degree specialism, indicating the prevalence of this subject and perceived importance of this subject within Computing degrees.

Although international universities were included within this review, 72% of the universities analysed were UK universities. 66% of the UK universities are part of the Russell Group, a “self-selecting body representing Britain’s foremost research-led universities” (The Observer, 2003). The remaining universities reviewed which were not part of the Russell Group were within the top 100 UK universities list (Complete University Guide, 2021) for example University of Bath (2021) is at number 11 and Northumbria University (2021) at number 52. It was important to review universities which were neither part of the top 10 or within the Russell Group to give a broader picture of the AI education landscape to determine if it is taught and how at varying institutions. The variation in entry requirements and pre-requisites

for the different universities enabled identification of the differing educational backgrounds of students and the educational levels at which this subject is taught.

Computing departments within the universities were the primary focus for the identification of modules within the umbrella of the AI domain. However, it is recognised that other departments may teach some form of AI due to the interdisciplinary nature of this domain, for example Machine Learning techniques may be taught within biological subjects to impart data analysis techniques to the students. It may also be possible that students from differing programmes may participate in modules run by the respective Computing departments which may raise issues with the more technical material due to differences in educational background.

4.3.2 Machine Learning Modules

Using the keywords “Machine Learning” when searching through an institution’s module catalogue enabled identification of the modules they offered within this specific domain. Out of the 25 universities reviewed, 22 offered some form of Machine Learning module. The modules identified were then initially categorised according to their title, 7 of the institutions offered modules entitled “Introduction to Machine Learning”, 10 provided “Machine Learning” modules and a further 9 offered modules which contained “Machine Learning” within the module title. Some of the universities under review offered more than one of these types of module, for example The University of Manchester (2021) listed 2 modules including “Foundations of Machine Learning” and “Machine Learning”. Both modules seemed to serve the same purpose of providing students with an introduction to this domain, however the content differed. The “Machine Learning” module also covered artificial neural networks and Deep Learning. Overall, 30 were offered across the 25 universities which included the term “Machine Learning”, 7 were “Introduction to Machine Learning” modules, 10 were “Machine Learning” modules and a further 13 were “Miscellaneous Machine Learning” modules. This excludes modules which were specifically listed as “Advanced Machine Learning.” It was important to determine at which level students were being introduced to this topic, therefore advanced courses were analysed separately.

4.3.2.1 Module Level

Determining at which educational level the Machine Learning modules were being offered would give an indication as to where universities place this domain within the curriculum as well as the complexity of material being delivered. Figure 9 shows the educational level at which the Machine Learning modules were being offered. The majority of these modules were offered at undergraduate level, such as the Machine Learning module at the University of Bath (2021) which aims to equip students with an understanding of the different Machine Learning algorithms and the application of these to real world data.

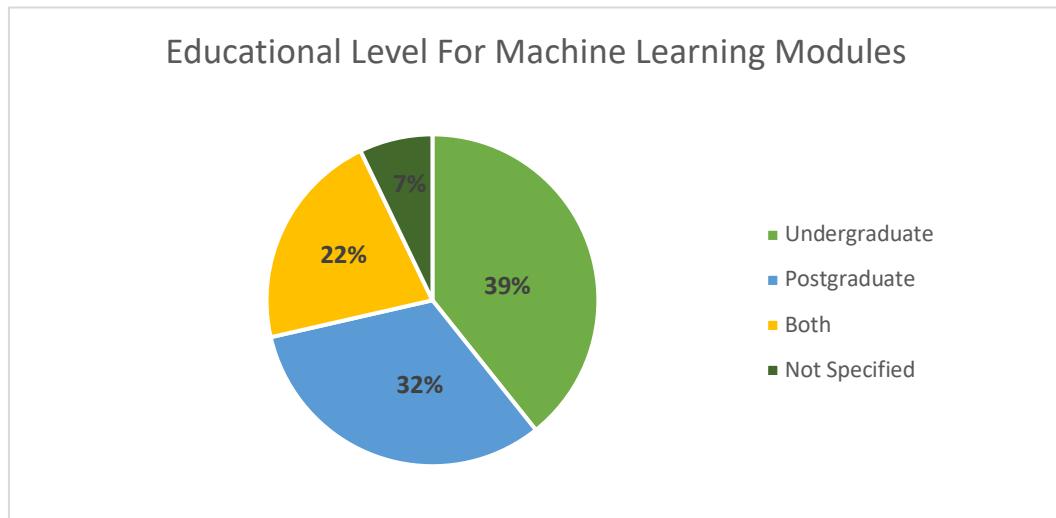


Figure 9: The educational level at which Machine Learning modules are offered

Within the online module catalogues it was often difficult to ascertain whether the module was optional or compulsory as the majority of the institutions did not list this information publicly. Therefore, it was not possible to be certain whether these types of modules were being offered within the degree programme as elective or mandatory.

4.3.2.2 Pre-requisites

Determining the pre-requisites for enrolment onto the Machine Learning modules identified in this review enabled insight into the expectations of the module leader. The type of pre-requisites also provided an indication of the material covered and at what difficulty level. Out of the 3 categories of Machine Learning modules reviewed, modules entitled "Machine Learning" had the highest number of pre-requisites. Out of the miscellaneous Machine Learning modules 46% of the modules specified pre-requisites and only 28% of the "Introduction to Machine Learning" modules had any prior requirements. Tables 4 and 5

display the pre-requisites for these Machine Learning modules, the table listings for the miscellaneous Machine Learning modules are included in Appendix G.

In all cases the pre-requisites were very similar in specifying prior knowledge and experience within specific aspects of mathematics such as calculus, linear algebra, probability and statistics, topics which are all pertinent to cognition of the theory of Machine Learning. The requirement of programming experience was also specified by some institutions including University of St Andrews (2021b) and MIT (2021). More formal pre-requisites were required by some of the institutions such as the University of Manchester (2021) who demand the students have completed prior modules in Mathematics and Data Science.

The analysis of pre-requisites did not take into account the requirements for entry onto the overall degree programme therefore, students may have a high level of mathematics or programming experience from previous learning. The natural variation of educational background amongst the cohort may cause issues when either teaching the math-heavy theoretical aspects of Machine Learning or the technical aspects, if students do not have extensive programming experience.

4.3.2.3 Content

The nature of content taught in these Machine Learning modules was difficult to analyse as some institutions were very vague with their descriptions and only outlined the types of algorithms taught such as supervised learning on their public webpages. However, other universities such as University of Manchester (2021) listed the specific models such as Support Vector Machine (Cortes and Vapnik, 1995). An overview of the content taught on each module is listed in Tables 4 and 5. The most common categories of Machine Learning algorithms included classification, regression, clustering and neural networks. This finding is fairly indicative of the Machine Learning domain (outlined in Figure 1, p.1) as classification and regression cover supervised learning and clustering is a form of unsupervised learning.

Where possible the categories of learning algorithms were drilled down further into specific models to gain a view of the details of the module. The most frequent algorithms mentioned were the Naïve Bayes Classifier (Rish, 2001), this was listed in 12 of the modules. The Support Vector Machine (Noble, 2006) and Decision Tree algorithms (Quinlan, 1986) were the next 2

most frequent. Six of the modules taught some form of Deep Learning, including Convolutional Neural Networks (LeCun *et al.*, 1998), Recurrent Neural Networks (Rumelhart and McClelland, 1987) and the backpropagation algorithm (Nielsen, 2015).

Introduction to Machine Learning Modules			
Institution	Educational Level	Pre-Requisites	Content
Massachusetts Institute of Technology (MIT)	Undergraduate	Calculus, Introductory Python programming module	Clustering, Classification, Probabilistic Modelling, SVM, Hidden Markov Models, Neural Networks
ETH Zurich	Postgraduate	None Listed	Regression, Classification, Neural Networks, Unsupervised Learning, Bayesian Approaches
Carnegie Mellon	Both Levels	None Listed	Classification, Regression, Unsupervised Learning, Bayes Rule
Berkeley	Not specified	None Listed	Classification, Regression, Neural Networks, Generative and Discriminative Models, Bayesian Networks, Clustering
Imperial College London	Both Levels	None Listed	Inductive Learning, Instance Based Learning, Hypothesis Evaluation, Neural Networks
University of Manchester	Postgraduate	Previous modules in: Mathematics, Data Science and AI	Classifiers, Linear Models, Algorithm Assessment, Feature Selection
University of Edinburgh	Both Levels	None Listed	Regression, Neural Networks, Clustering, Margin Based Methods

Table 4: Overview of the Introduction to Machine Learning modules

The content between the educational levels was very similar, however there were some indications that the postgraduate modules covered more in-depth topics as would be

expected. For example, Principles of Machine Learning Systems (University of Cambridge, 2020b) covers fairly advanced topics such as an exploration of Deep Learning compilers.

Machine Learning Modules			
Institution	Educational Level	Pre-Requisites	Content
Stanford	Both Levels	Linear Algebra, Probability, Statistics	Basic concepts, generative learning, evaluating and debugging learning algorithms
National University of Singapore	Undergraduate	Previous modules	Naïve Bayes, Linear Classifiers, Neural Networks, Deep Learning
Oxford University	Both Levels	Probability, Linear Algebra, Calculus, programming skills	Linear Prediction, Bayesian Machine Learning, Neural Networks, Clustering
University of St Andrews	Postgraduate	Previous module: Object-Oriented Programming, A-Level Maths	Mathematical foundations, Regression, Classification, unsupervised learning
University of Manchester	Undergraduate	None listed	Neural networks, SVM, Clustering, Deep Learning
University of Bath	Undergraduate	None listed	Unsupervised learning, Classification, Regression
University of Glasgow	Both Levels	Knowledge of maths (e.g., matrices, linear spaces, geometry)	Regression, Classification, Clustering, Probability Density Estimation, Dimensionality Reduction
University of Leeds	Undergraduate	Previous module: Artificial Intelligence	Neural Networks, Bayesian Learning, Clustering, Reinforcement Learning, Deep Learning
University of Bristol	Not Listed	Linear Algebra, Multivariate Calculus	How to build models of data
Newcastle University	Postgraduate	Non listed	Supervised and unsupervised learning, Deep Learning

Table 5: Overview of the Machine Learning modules

Only 10% of the modules analysed taught some form of ethics, or the legal and social issues surrounding the use and application of Machine Learning. This lack of inclusion raises concerns around the current education provision relating to the differing potential impacts of this

technology and the need for greater inclusion to ensure students completing these courses are equipped with the knowledge to become responsible practitioners.

4.3.2.4 Module Structure

Establishing how the modules were structured was challenging as this information was often not listed online, or accessible without an institution login. From the limited information available there was a mix of one semester and two semester courses. The Machine Learning modules also often followed the convention within Computing courses of teaching material through a mixture of lectures and practical sessions. Practical sessions allow the students to embed their theoretical knowledge within a practical scenario within a supportive environment.

4.3.2.5 Assessment

Out of the 30 modules analysed, assessment methodology could only be found for 18. The most popular type of assessment for these modules, at 61%, was a mixture of coursework and some form of exam. The weightings of the types of assessment differed per module, however in each case the exam carried a higher percentage weighting for the overall assessment. This mix of assessment is often used within Computing courses to evaluate the student's theoretical knowledge as well as their practical application skills (Barnes *et al.*, 2020). 33% of the modules offered coursework only assessment, these modules tended to have a higher level of practical content, for example University of Edinburgh's (2018) Machine Learning Practical which is focused on the implementation and evaluation of Machine Learning systems.

4.3.2.6 Learning Outcomes

Where available, the learning outcomes were assessed to determine the aims and outcomes for Machine Learning modules. The learning outcomes followed a similar pattern in that students were expected to be able to explain the different Machine Learning techniques covered within the syllabus as well as discuss the scope and limits to each method. Upon completion of the module the students were expected to be able to apply the various Machine Learning methods to given datasets. One of the other main goals specified within the module

descriptors was the ability to differentiate and evaluate the different approaches as well as appreciate the limitations and capabilities of Machine Learning.

Advanced Machine Learning Modules

Four out of the twenty-five universities reviewed offered a specific Advanced Machine Learning module. As shown in Table 6, 3 out of the 4 institutions required the students to have knowledge of a range of mathematical concepts, particularly in statistics, linear algebra and calculus. Established programming skills were also essential for 2 of the modules. Although prior experience in a particular language was not specified, both of these modules used Python (Python Software Foundation, 2021) as the main programming language.

Advanced Machine Learning Modules			
Institution	Pre-Requisites	Content	Assessment
ETH Zurich	Analysis, statistics, numerical methods for CS, programming skills	Bayesian learning, computational learning theory, non-parametric density estimation	Exam and coursework
Oxford University	Probability theory, linear algebra, continuous mathematics, multivariate calculus, programming skills	Mathematics for ML, Bayesian modelling, Gaussian processes, GPU optimization for deep neural networks	Not Listed
University College London	Linear algebra, probability theory, calculus	Kernel methods, ICA, SVM, regularization	Exam and coursework
Carnegie Mellon	Not Listed	Non-parametric Bayes, high-dimensional regression, deep density estimation, decision theory	Coursework

Table 6: Overview of Advanced Machine Learning modules

The content on these more advanced courses differed from the previous Machine Learning modules as they covered more theoretical aspects of this domain such as computational learning theory and mathematics for Machine Learning. However, there were similarities in that some form of Bayesian learning was taught on the more introductory modules and was also covered on these advanced modules. The assessment strategies for these types of modules were similar to the Machine Learning modules in that a mixture of an exam and coursework was the most popular method.

4.3.3 Deep Learning Modules

Nine of the universities reviewed as part of this data collection method offered some form of Deep learning module. Overall, 10 modules were analysed which included the key words “Deep Learning.” University College London offered two Deep Learning modules, one introductory course and one more advanced (University College London, 2021). One observation of this analysis process was that information pertaining to these modules was harder to find than the Machine Learning modules.

Information pertaining to the educational level of the module could not be found for 4 of the modules. However, 2 of the Deep Learning modules were offered at undergraduate level, 3 at postgraduate and 1 was offered to students at both levels. Seven out of the ten modules specified pre-requisites. The pre-requisites were similar to the Machine Learning modules in that students needed to be familiar with specific areas of mathematics and statistics as well as programming experience. However, the pre-requisites were extended for Deep Learning modules in that students had to have passed previous modules in Machine Learning, such as at Durham University (2021) where students had to have done Artificial Intelligence and Data Science modules.

The most frequent Deep Learning models taught within these modules included the Convolutional Neural Network (LeCun *et al.*, 1998), Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997) and Generative Adversarial Networks (Goodfellow *et al.*, 2020). Four of the modules also covered content relating to reinforcement learning. Alongside the specific Deep Learning algorithms, regularization, optimization and the backpropagation algorithm were taught. The most popular form of assessment for the Deep Learning coursework was purely coursework with 40%, perhaps reflecting the more applied nature of these types of modules.

4.3.4 Artificial Intelligence Modules

Out of the 25 universities under review, 15 offered some form of Artificial Intelligence module. Overall, there were 19 modules reviewed which contained the keyword ‘Artificial Intelligence.’ Three of the universities offered more than one module, for example the University of St Andrews (2021a) offered one undergraduate module entitled Artificial Intelligence, which

covers the general features of AI. They also host two postgraduate AI modules, one covering AI practice which instructs on the practical design of this topic and another on AI principles. Ten of the AI modules were offered at undergraduate level, 3 were postgraduate and 2 were offered on both undergraduate and postgraduate courses. Out of the universities analysed, 4 did not list the educational level at which their AI module was offered.

Thirteen out of the AI nineteen modules had specified pre-requisites, 4 modules did not appear to require any prior knowledge and 2 did not list this information. Out of all of the educational levels, the undergraduate modules required the most pre-requisites. Nine of the modules required more formal pre-requisites in the shape of prior modules being undertaken. Eight out of the nine modules required prior completion of some form of mathematics and programming module. Other domain specific requirements were listed, for example the postgraduate Artificial Intelligence Practice module at University of St Andrews (2021a) required the student to have already completed an AI module. At the University of Manchester (2021) students were required to have already completed a module in Data Science before they could start the Introduction to AI module. Out of the universities which had fewer formal pre-requisites, students were required to have knowledge and experience in specific aspects of mathematics including probability, linear algebra and basic geometry and programming experience.

Module content varied as shown in Table 7, which lists the content for the undergraduate AI modules. However, amongst all of the AI modules offered at a variety of educational levels, nearly all of the modules offered content related to search, some form of game playing and logical reasoning. Thirteen of all of the AI modules analysed taught some form of Machine Learning, including neural nets and Deep Learning. Philosophical issues and the ethics of AI were also included within the content for many of these modules.

Determining the assessment methods for these modules was often difficult, similar to the Deep Learning modules. However, for the modules which did list assessment techniques, the most popular form of assessment was a mixture of coursework and exam with 9 modules employing this method. Six of the modules graded their students solely through coursework.

Undergraduate Artificial Intelligence Modules		
Institution	Module Title	Content
Massachusetts Institute of Technology	Artificial Intelligence	Rule chaining, constraint propagation, constrained search, Support Vector Machines, neural nets, generic algorithms
National University of Singapore	Introduction to Artificial Intelligence	Game playing, logic, uncertainty, probabilistic reasoning, Machine Learning
University of St Andrews	Artificial Intelligence	Search, games, reasoning about uncertainty, Machine Learning, philosophy of AI
Durham University	Artificial Intelligence	AI search, ethics and bias in AI, Machine Learning
University of Manchester	Introduction to AI	Search and planning, logic and reasoning, AI and probability, knowledge representation, philosophical issues
University of Bath	Artificial Intelligence	Problem solving through state-space search; logical reasoning; probabilistic reasoning; Machine Learning, social, legal, and ethical implications of AI
University of Leeds	Artificial Intelligence	Search techniques, logic, knowledge representation, Markov models, ethical issues, game play and searching
University of Aberdeen	Grand Challenges of Computing and Artificial Intelligence	NLP, computer vision, robotics, search, neural networks, reinforcement learning
University of Sunderland	Artificial Intelligence	Knowledge representation and reasoning, search, introduction to Machine Learning
Northumbria University	Artificial Intelligence and Robotics	Machine Learning and Deep Learning, evolutionary and genetic algorithms

Table 7: Overview of undergraduate Artificial Intelligence modules

4.3.5 Data Science Modules

Out of the universities analysed 2 offered a specific Data Science module at the time of data collection, these universities were University of Cambridge (2020a) and University of Sunderland (2021). Both of these modules were at postgraduate level, one was titled Data Science: Principles and Practice (Cambridge), the other Data Science Fundamentals (Sunderland). The module hosted by the University of Cambridge required previous modules to be undertaken in Mathematics and Machine Learning, however the University of Sunderland module had no prior pre-requisites.

The content was similar for both modules, topics included linear regression and some form of classification. However, the University of Cambridge module covered Deep Learning whereas

the University of Sunderland module did not. Assessment for both of the modules was purely coursework based.

4.3.6 Summary of Findings

All of the 25 universities reviewed offered some form of module which could be categorised within the domain of Artificial Intelligence, whether this was a Machine Learning or Deep Learning module. This prevalence of education provision may suggest a more permanent place for these types of modules within the Computer Science curriculum. However, the variation in modules offered and disparity in content covered on modules similarly titled suggests that there is no real agreement on the core pedagogical content needed for such modules.

Five of the universities offered separate modules in AI, Machine Learning and Deep Learning. Within the content for some of the modules, particularly the ones which contained “Machine Learning” in their title, content was included which covered the general principles of AI, Machine Learning and Deep Learning. There was also overlap in content with many Machine Learning modules covering Deep Learning, which is to be expected as Deep Learning is a sub-domain of Machine Learning. However, some of the universities chose to separate these fields into separate modules, perhaps due to the wealth of topics which can be covered within these subject areas. The universities also offered specialisations within the field of AI, many ran courses on computer vision, natural language processing or robotics. However, this was out of the scope of the outline for this particular round of data collection. Although, it is worth noting the wealth and variety of modules now being offered within this domain.

Out of the categories of modules which were analysed (“Machine Learning”, “Advanced Machine Learning”, “Deep Learning”, “Data Science”) these modules were offered at a variety of educational levels. The Machine Learning and Artificial Intelligence modules were mainly offered on undergraduate programmes, the Deep Learning and Data Science modules were principally postgraduate modules. An overarching educational level could not be determined for the Advanced Machine Learning modules.

The Deep Learning modules had the highest number of modules out of the different types analysed with outlined pre-requisites, 70% (7) of the Deep Learning modules had prior

requirements. 40% (4) of these requirements were relating to programming skills and 50% (5) related to mathematics knowledge, particularly in calculus, linear algebra, probability and statistics. 50% (5) of the Deep Learning modules also required knowledge in Machine Learning, usually through a prior module. Artificial Intelligence modules were the next highest in regard to module pre-requisites at 68% (13) for all of the courses analysed including this key word. 50% of all 30 of the modules including the keyword “Machine learning” (excluding advanced courses) had prior requirements. Twelve of the modules required prior knowledge of mathematics topics related to the field, 7 of the modules stated students should have programming experience and 6 of the modules required both programming and mathematics knowledge.

One of the common themes which emerged in relation to module pre-requisites was the requirement for both programming experience and knowledge within particular areas of mathematics. The majority of modules which listed mathematics as a core requirement often specified knowledge was needed in relation to linear algebra, calculus, probability and statistics.

As previously mentioned, there was some commonality amongst the different types of modules analysed regarding course content which was to be expected due to the nature of the field of AI. However, alongside the commonalities, it was also important to distinguish any differences in content specific to the type of module. The prominence of Machine Learning amongst the differing modules was noteworthy, with many covering classification, regression and some form of neural networks. The most common models/algorithms included Naïve Bayes Classifier (Gandhi, 2018), Support Vector Machine (Cortes and Vapnik, 1995), Decision Trees (Quinlan, 1986), Convolutional Neural Network (LeCun *et al.*, 1998) and the Recurrent Neural Network (Hochreiter and Urgen Schmidhuber, 1997).

The type of assessment for the module seemed to reflect the nature of the content, for example the Data Science modules which were assessed purely on coursework, where these courses had a more applied nature. This was similar for the Deep Learning modules, where the most popular form of assessment was coursework based. However, the Machine Learning and AI modules used a mixture of coursework and exam for assessment, reflecting an importance on the theoretical as well as practical elements of the content.

4.4 Data on Lecturers' Experiences

Two types of data were analysed relating to lecturer's experience teaching within this domain, these included questionnaires and interviews. These methods were employed to gather further data on how these types of modules are currently being taught, as well as to gain insight from lecturers on their experiences teaching within this domain. There was a low response rate for both of the data collection methods for lecturers. This response rate was to be expected as unsolicited emails were sent calling for participation in the study through either questionnaire or interview.

4.4.1 Questionnaires

The low response rate, of 3 respondents, to the questionnaire for lecturers teaching AI means that the results are limited, however the questionnaire responses do provide a good indication of the expectations the lecturers have of the students, the type of content they teach and specific teaching strategies they employ. Two of the three lecturers who responded taught their module at postgraduate level as an optional choice. In relation to demographics, the respondents indicated that their gender demographics were in line with other computing modules in that the majority of students were male. An expectation of some form of prior mathematics knowledge was outlined, particularly in statistics, linear algebra and probability. However, the majority of module leaders expected the students to be beginners with programming, particularly with the Python programming language.

The most popular topics listed as part of the module content included Linear and Logistic Regression (Worster, Fan and Ismaila, 2007), K-Means (Bock, 2008), Principal Component Analysis (Jolliffe and Cadima, 2016), Support Vector Machine (Cortes and Vapnik, 1995), Bayesian Machine Learning (Algorithmia, 2020) and topics related to neural networks. All of the respondents stated that they taught some form of Deep Learning including Convolutional Neural Networks (LeCun *et al.*, 1998), Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997) and the backpropagation algorithm (Rumelhart, Hinton and Williams, 1986). The most prominent form of assessment from the frequency analysis was the combination of coursework and exam. The lecturers were asked to list which additional resources they advise students to use and where to look for further information, online resources were the most popular response.

Lecturers were asked, in their experience where they felt students encountered the most difficulties or where they struggled. Two out of the three responses mentioned mathematics. One of the respondents mentioned the modelling experience and advised that the “*students need to do projects to master it.*” Another respondent mentioned that “*more lab sessions would be helpful.*” The questionnaire also enquired into student feedback on the module, all of the respondents stated that their feedback was positive. However, two of the respondents made additional comments, one stating that students “*claim the pace was fast*” and another stating that “*some students complain about the amount of material.*” Within the additional comments section of the questionnaire one of the respondents advised that the students get easily overwhelmed in the lectures and that “*students especially with a background on stats/math find it much easier than the ones without technical background.*”

4.4.2 Interviews

Overall, there were 5 respondents to the call for participation in the interviews. Four of the interviews were completed in person and one via online call. Through thematic analysis, the following descriptive units were uncovered, these were “prerequisites”, “pedagogy”, “content”, and “perceived difficulties.” Inter-rater reliability analysis was completed to determine the degree of agreement between the two raters in relation to the thematic analysis categories. The interrater reliability for the raters was found to be Kappa = 0.48 which can be interpreted as a moderate agreement (McHugh, 2012).

4 out of the 5 instructors taught courses delivered as part of an undergraduate programme.

The module titles the participants taught included:

- Machine Learning, Computer Vision (two modules)
- Artificial Intelligence
- Machine Learning and Computer Vision
- Machine Learning and Natural Language Processing
- Data Mining and Machine Learning

One of the participants taught two separate modules within this domain, a general Machine Learning module as well as a computer vision course. Computer vision was also taught within one of the other modules from another respondent. Alongside computer vision, natural

language processing was also a specialism included within another Machine Learning module. The Artificial Intelligence module was taught as an introductory module and covered all content within the AI domain, including Machine Learning.

4.4.2.1 Pre-Requisites

Four out of the six modules discussed during the interviews had pre-requisites. Three of these were formal requirements pertaining to previous modules which the students had to complete before starting. These prior modules were related to programming and mathematics and one required a previous Machine Learning module to be completed. The module, which did not have formal pre-requisites required the students to have prior programming experience.

4.4.2.2 Content

As part of the interview, lecturers were asked two questions relating to the content of the modules they teach. They were firstly asked what topics they taught as part of their module as shown in table 8. Participants were also asked which topics within this domain they consider pivotal to teach.

Topics Mentioned	
Neural Networks	Backpropagation
Bayesian methods for Machine Learning	Recurrent Neural Networks (RNN)
Supervised Learning	Convolutional Neural Networks (CNN)
Clustering	Unsupervised Learning
Support Vector Machine (SVM)	Classification
Gaussian Mixture Models (GMM)	Decision Trees
Ethical Issues	Introduction to Statistics

Table 8: List of all topics mentioned by participants within the interview

When discussing the content on their modules, some of the lecturers mentioned specific Machine Learning algorithms such as the Support Vector Machine (Cortes and Vapnik, 1995), whereas other participants were more general mentioning that they taught supervised and unsupervised learning. There was also overlap of content in some responses in that they mentioned that they taught neural networks and then listed specific types of networks such as RNNs (Hochreiter and Urgen Schmidhuber, 1997) and CNNs (LeCun *et al.*, 1998). Introduction to statistics was taught on one of the modules which had no pre-requisites, to

ensure the students had a level of mathematics knowledge in order to fully understand the other content on the module.

The module leaders were also asked which specific topics they consider pivotal to comprehend within this domain. Interestingly, the most popular response was Deep Learning and specific models such as CNNs and LSTM (long short-term memory) (Brownlee, 2017). Other responses included association rule learning and clustering.

4.4.2.3 Pedagogy

Questions relating to pedagogy included specific strategies to help aid student learning, additional information and resources provided to students and assessment techniques. All of the modules were taught through a mixture of lectures and practical sessions to enable students to comprehend the theory and then undertake practical examples. One of the interviewees advised that they *“think students learn better by doing something”* and that they *“try to make it easy to see an end result”* for the practical task they set for students. One of the participants also ran an “AI Lab” for both staff and students if they required additional support. Real-life scenarios were also mentioned as a tool to help facilitate student learning, as these not only motivated the students but helped them to contextualise the theoretical knowledge they were learning. One of the lecturers used online quizzes and an online forum to facilitate learning and provide support to the students.

Four out of the five lecturers interviewed referred students to online resources as additional material to help aid their learning. Amongst the online resources, two of the lecturers specified that they refer students to specific online videos on YouTube to supplement the material covered within their lectures. The use of textbooks and research papers were also mentioned as learning resources. Three of the lecturers used a combination of assessment techniques, consisting of an exam and practical coursework. The remaining two purely assessed using coursework.

4.4.2.4 Perceived Difficulties

One of the pivotal questions within the interview was determining, based upon the participants experience teaching this subject, what they felt students struggled the most with.

This would help determine any barriers students are encountering when undertaking courses within this area. It was also recognised that this may be a sensitive topic to discuss as participants may feel any answer they give may be a reflection on their teaching skills and may have been reluctant to fully answer this question. However, the majority of the participants were open and forthcoming in answering this question.

Three of the five participants mentioned students had difficulty with the theoretical aspects of the module, two of these specifically mentioned issues pertaining to mathematics, one participant stated, “*especially equations.*” Other issues identified included the terminology used within this field and that students struggled to match the dataset to an appropriate algorithm. One of the participants also shared that they have received feedback from students that “*they found it too hard.*” One of the participants stated that their students did not struggle with any aspects of their module, however they added the caveat that students who were not particularly confident with mathematics would not choose to take the module they ran.

4.4.3 Key Themes

From the data collected gathering the opinions of individuals who are responsible for teaching courses within AI, it was apparent that these types of modules are offered to students at both undergraduate and postgraduate level. Similar to other modules within Computing Science, the content is taught through a mixture of lectures and practical sessions. However, modules within the AI domain, often have a number of pre-requisites for admittance onto the module, most commonly a high level of mathematics knowledge.

As would be expected, the majority of modules cover both supervised and unsupervised learning. Common algorithms included the Support Vector Machine (Cortes and Vapnik, 1995), Bayesian Machine Learning techniques (Algorithmia, 2020) and neural networks. Convolutional Neural Networks (LeCun *et al.*, 1998), Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997) and the backpropagation algorithm (Rumelhart, Hinton and Williams, 1986) were commonly taught. Students were often pointed to online resources for supplementary material to aid their learning. The most common form of assessment was a mixture of exam and practical coursework.

From the accounts of the lecturers, it is apparent that they feel students struggle the most with the theoretical aspects of this field, particularly with the mathematics. Even though the majority of the modules require students to have prior knowledge of mathematics, this is still an issue for some students. This suggests a wider problem which needs to be addressed to help students overcome difficulties when learning this field. The first-hand accounts conveyed by the lecturers, within the questionnaires and interviews relating to students feeding back to them that they have found the module too hard, or that the pace is too fast and there is too much material suggests an underlying issue with module structure and content, essentially the pedagogical content knowledge.

4.5 Case Studies

The results from the three case studies are presented by institution below, these sections contain analysis of the differing collection methods employed. The key themes will be summarised in the section following the institutional reviews.

4.5.1 Newcastle University Machine Learning Module

The Machine Learning module at Newcastle University (University A) is offered at postgraduate level and is a requirement for students from two different course streams. The students are either enrolled on a master's programme with a specialisation in Data Science or as part of a CDT (Centre for Doctoral Training) which has a focus on Cloud Computing. Relating to the educational backgrounds of these students, there is some disparity. The CDT requires a higher level of prerequisite mathematics knowledge than the Data Science programme, however, both specify the same entry qualification for an undergraduate degree in a relevant discipline (e.g., Computing or Engineering).

The module is block taught lasting three weeks. This format has been suggested to "improve engagement, attendance and attainment, particularly amongst students from diverse entry pathways" (Dixon and O'Gorman, 2020). The first two weeks of the module consist of the main taught content. The content is taught through a mixture of lectures and practical sessions where students apply the theoretical knowledge they have gained, within a supported environment. In the third week of the module, the focus turns to how this technology is applied in a number of differing domains and for varying applications such as computer vision.

This content is delivered through guest lectures from industry professionals working in these areas. The assessment pathway is the same for both cohorts of students in that it is coursework based.

4.5.1.1 Pre-Module Questionnaire

Completion of the pre-module questionnaire was optional and was offered for completion to two student cohorts. Out of the first cohort, 21 students completed the questionnaire, this was a response rate of approximately 35% from the average module attendance of 60. Out of the second cohort, 35 students participated, with an approximate response rate of 58% (out of 60). In total 56 students completed this questionnaire. These relatively high response rates could be a consequence of the method for participation. The students were approached during their practical session with a paper-based form to complete.

Student Demographics

Determining the student demographics was important to comprehend who was studying these types of modules. The data collected in relation to demographics included age and gender. The majority of students at 32%, were within the age range of 23-27 years old as displayed in Table 9, closely followed by the 18-22 age range.

		Age		
		Frequency	Percent	Valid Percent
Valid	1. 18-22	17	30.4	30.4
	2. 23-27	18	32.1	32.1
	3. 28-32	7	12.5	12.5
	4. 33-40	9	16.1	16.1
	5. 40+	3	5.4	5.4
	6. Prefer not to say	2	3.6	3.6
	Total	56	100.0	100.0

Table 9: Age demographics for students on Machine Learning module

As shown in Table 10, 71% of the students enrolled on the Machine Learning module identified as male when asked about their gender and 25% identified as female. One student preferred to self-describe and another preferred not to disclose this information. The gender demographics of students for this module match the national picture for underrepresentation of women in Computer Science. In the academic years in which this questionnaire was

completed, only 15% of graduates in Computer Science were women in 2017/18 and 16% in 2018/19 (STEM Women, 2021).

		Gender			
		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	1. Male	40	71.4	71.4	71.4
	2. Female	14	25.0	25.0	96.4
	3. Prefer not to say	1	1.8	1.8	98.2
	4. Prefer to self-describe	1	1.8	1.8	100.0
	Total	56	100.0	100.0	

Table 10: Gender for students on the Machine Learning module.

Educational Background

A number of questions were included within the questionnaire to understand the educational background of students undertaking modules within this domain, including questions pertaining to mathematics and programming experience and confidence. As this was a postgraduate module, students were asked to list their undergraduate degree. The most frequent response, as shown in Figure 10, was Computing at 39%, followed by Mathematics. Surprisingly, two of the questionnaire respondents had a PhD, perhaps reflecting the trend for re-skilling in Data Science areas. The degrees listed within the “other” category included topics such as Finance, Materials Science and Chemistry.

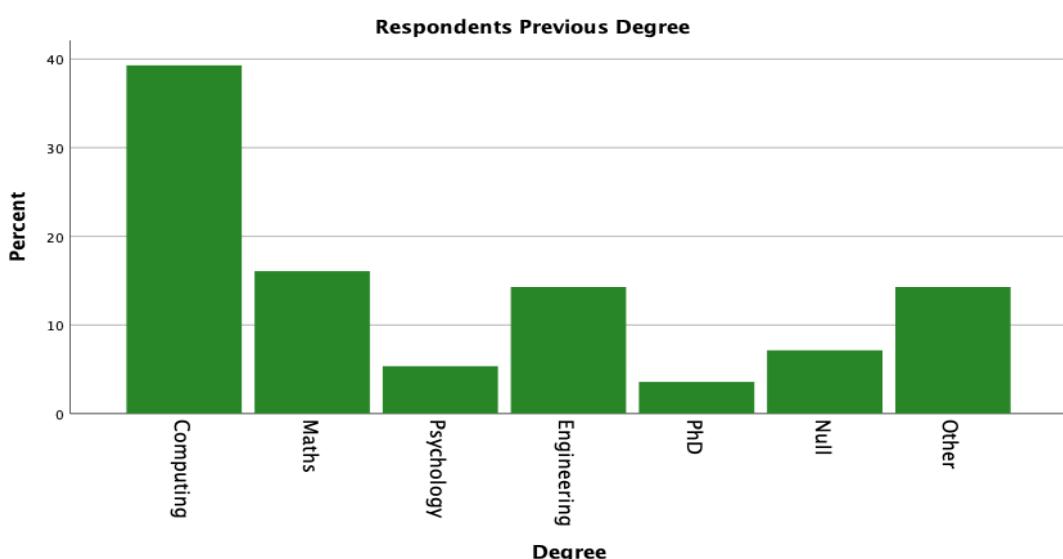


Figure 10: Previous degree of students from Newcastle University

The participants were asked what their highest level of mathematics attainment was, the majority of students (37%) had A-Level Mathematics, 26% stated that mathematics was a

major part of their first degree. However, some of the students (16%) had a mathematics attainment level of GCSE or equivalent. Alongside mathematics attainment, the students were also asked to rate their confidence in mathematics on a scale from 1 to 10 (10 being exceptionally confident). The mean confidence level was 6.6.

Before undertaking any inferential statistical tests, the data was examined to determine normality as this would influence whether parametric or nonparametric tests were most suitable. One-way between groups ANOVA (refer to Section 3.4.1, p.77) was carried out to determine whether mathematics attainment had an impact on self-reported confidence level (Table 11). Determining this could potentially indicate whether a lack of mathematical educational background is a barrier to learning (RQ1) and whether mathematics anxiety is potentially a difficulty encountered by students (RQ2). There was a statistically significant difference at the $p < .005$ (Sig .000 means $p < .005$) for the five groups of mathematics attainment. The effect size, calculated using eta squared, was 0.49, indicating a large effect size (Geert van den Berg, 2021). Post-hoc comparisons using the Tukey HSD test (Beck, 2018) indicated that the mean score for A-Level ($M = 6.67, SD = 1.238$) was significantly different from GCSE ($M = 4.44, SD = 1.509$). The mean score for the category: major part of degree ($M = 7.13, SD = 1.125$) was significantly different from GCSE. The mean score for the first degree group ($M = 8.22, SD = 1.093$) was significantly different from GCSE, A-Level and the respondents who chose not to disclose their mathematics attainment. These results suggest, as would perhaps be expected, that the higher the mathematics attainment level, the higher the self-reported confidence level.

ANOVA

Confidence in Maths

	Sum of Squares	df	Mean Square	F	Sig.
Between Groups	74.947	4	18.737	12.544	.000
Within Groups	76.178	51	1.494		
Total	151.125	55			

Table 11: Confidence in Maths ANOVA

Alongside determining students' mathematics skills, students were also asked to self-rate their programming skills as either "beginner", "novice", or "expert." 60% of students described themselves as a "novice" programmer, 23% as an "expert" and 14% as a "beginner". Similar to the mathematics section of the questionnaire, the students were asked to rate their

confidence in their programming skills. The mean for programming confidence was 6.2 out of 10. Table 12 displays the spread of self-determined programming skill level by gender, this shows that although the majority of male and female students identify as novice programmers at 60%, more women than men identified themselves as beginners.

Gender * Level of programming Crosstabulation

			Level of programming				Total	
			1 Beginner	2 Novice	3 Expert	4 No answer		
Gender	1 Male	Count	3	25	11	1	40	
		% within Gender	7.5%	62.5%	27.5%	2.5%	100.0%	
		Adjusted Residual	-2.3	.4	1.2	.6		
	2 Female	Count	5	8	1	0	14	
		% within Gender	35.7%	57.1%	7.1%	0.0%	100.0%	
		Adjusted Residual	2.6	-.3	-1.6	-.6		
	3 Prefer not to say	Count	0	0	1	0	1	
		% within Gender	0.0%	0.0%	100.0%	0.0%	100.0%	
		Adjusted Residual	-.4	-1.3	1.8	-.1		
	4 Prefer to self-describe	Count	0	1	0	0	1	
		% within Gender	0.0%	100.0%	0.0%	0.0%	100.0%	
		Adjusted Residual	-.4	.8	-.6	-.1		
Total		Count	8	34	13	1	56	
		% within Gender	14.3%	60.7%	23.2%	1.8%	100.0%	

Table 12: Gender and programming skill level crosstabulation

To further explore whether there was an association between gender and the self-rated programming skill level, the chi-square test for independence was used (refer to Section 3.4.1, p.78). However, the minimum expected cell frequency was violated due to the sample sizes. Instead, Fishers Exact Test (Sprent, 2014) was used as this is very similar to the chi-square test but is appropriate for small sample sizes, there was no statistical significance found as the value returned was 34.984 resulting in a p value of .114.

Prior Knowledge of Machine Learning

Students were asked to write a brief description of what they thought Machine Learning is to determine any mental models or preconceptions students might have when starting a module within this area. The responses to this question underwent inductive thematic analysis and were coded into eight different categories. Table 13 shows the categories and frequencies for these categories. The most frequent response was related to learning/extracting knowledge from data. Seven of the participant's responses could not be categorised or thematically grouped under an umbrella term like the other responses due to their uniqueness of response, therefore these were grouped within the "other" category. For example, one respondent wrote "*it is used quite widely by a lot of companies*", another student stated that they weren't sure and were here to learn.

Description of Machine Learning

		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	1 Learning/extracting knowledge from data	13	23.2	23.2	23.2
	2 No answer	7	12.5	12.5	35.7
	3 Automation/self program	8	14.3	14.3	50.0
	4 Modelling	3	5.4	5.4	55.4
	5 Pattern recognition	5	8.9	8.9	64.3
	6 Prediction	9	16.1	16.1	80.4
	7 Classify data	4	7.1	7.1	87.5
	8 Other	7	12.5	12.5	100.0
	Total	56	100.0	100.0	

Table 13: Responses to the description of Machine Learning

As the Machine Learning module was compulsory, the students were asked to rate how interested they were in learning this topic on a scale of 1 to 10 (10 = very interested). Motivation is linked with self-regulatory processes, particularly "want-to" motivation which is also associated with better goal attainment (Werner and Milyavskaya, 2019). The most frequent response at 30% was 9 out of 10, closely followed by 10 at 26%. None of the respondents rated their motivation below a 5.

As well as asking the students to rate their confidence in mathematics and programming, again on scale of 1-10, they were also asked to rate their confidence in their ability to do well

overall within the module. The results for this question were mixed, with some students rating their confidence as low as 2 and 3. The most frequent responses were 5 and 6 at 23% and 21% respectively. The students were also asked to assess which aspect of the module they expected to find the most difficult, either the theory or the practical aspects. 71% of the students stated that they expected to find the theoretical side of the module more difficult than the practical side. Relating to demographics there was no relationship present between gender and which aspect of the module they expected to find most difficult.

Teaching and Learning Strategies

A number of questions were included within the questionnaire to determine which teaching and learning strategies/methods students perceive as the most useful and which they intended to engage with for their learning within this module. They were first asked which skills from their previous learning they would use for this module. The most popular responses were motivation with 55% (31 respondents) stating 'yes' they would use this, programming skills with 76.8% (43), numeracy skills with 57% (32) and analytical skills with 75% (42). Interestingly this cohort of students indicated that they would not be using more holistic strategies for learning such as perseverance and resilience.

The participants were asked to rate a number of learning tools/methods including lectures, handouts, quizzes, practicals and assignments on a scale from 1 to 5 (5 = not important) in terms of importance to their learning. Out of all of the methods listed, lectures had the highest importance score with 25% (14) of the students giving it a 1 out of 5 on the scale. The results relating to practical sessions was interesting as 26% (15) of the students voted this as a 5 on the scale of importance, therefore stating that these were not important, however, 25% (14) of students voted this as a 1 on the scale of importance. Therefore, students were split between finding practical sessions important and not important. Spearman's rank order correlation (refer to Section 3.4.1, p.76) was used to determine if there was a relationship between students who found the practical sessions important and how they ranked their programming skill level. Spearman correlation determines the strength and direction of a monotonic relationship, where the variables tend to change together, but not necessarily at a constant rate (Minitab, 2019). The correlation coefficient expresses the strength of the relationship between the two variables. The correlation coefficient for our test was .056. This

indicates that there was no significant relationship between students programming skill level and how important they ranked practical sessions.

Students were asked which additional resources they would use for learning, 57% of students chose not to answer this question, perhaps as it was not multiple choice and required students to write and consider their own options. However, handouts which summarise the material and practical examples were the most popular responses for participants who chose to complete this question. Participants were asked to rate on a scale of 1 to 6 (1 = first choice), where they would prioritise going for advice or support. Out of all of the options listed (refer to Appendix C for question options), online resources were most often voted the first choice at 39%, textbooks were most often voted as the last choice at 44%.

Finally, students were asked what learning strategies they proposed to use within this module. Table 14 shows the responses to each strategy. The strategies highlighted in grey are the ones which students acknowledged they would use within the Machine Learning module. These included note taking, practical exercises, online guidance and critical thinking.

Learning Strategies (%)			
Learning Strategy	Yes- they will use	No – they will not use	No answer
Note Taking	57.1	41.1	1.8
Study Groups	16.1	82.1	1.8
Practical Exercises	91.1	7.1	1.8
Quizzes	19.6	78.6	1.8
Textbooks	35.7	62.5	1.8
Online Guidance	67.9	30.4	1.8
Critical Thinking	51.8	46.4	1.8
Reflection	33.9	64.3	1.8
Goal Setting	30.4	67.9	1.8
Planning	44.6	53.6	1.8
Self-Evaluation	28.6	69.6	1.8

Table 14: Learning strategies students propose to use in module

The findings presented in Table 14 need to be considered within the educational context and consideration given to the influence a lecturer provided resource may have on these results. For example, practical exercises provided by the lecturer may be deemed a valuable learning strategy as they have been provided by the person leading on the module and therefore deemed more important and relevant to building skills for the assessment. It may be valuable

to provide activities which embed a variety of learning strategies to actively promote the ones which students do not deem particularly useful but which may assist in their learning such as self-evaluation and reflection (RQ3).

4.5.1.2 Observation

Observation was undertaken with the 2018/19 cohort of students at Newcastle University (University A). As a standardised observation guide (discussed in Section 3.3.2, p.70) was used during the observation sessions, this aided the content analysis process. The data was firstly transcribed, then categories were created based upon the observation guide (appendix D). The content was then coded and placed into the categories. These categories include lecture context, topics, lecturer/student interaction, student to student interaction, pedagogy and then general comments from the observation session.

Attendance levels for these lectures were usually around 60 students per session, however attendance dropped for the last week of the module which consisted of guest lectures. A wealth of material was covered throughout the three weeks of the module, within the first lecture, the lecturer stated that this module is “*very challenging*.” Topics included supervised and unsupervised learning, Deep Learning and specific domain applications of these technologies. Specific models included the Support Vector Machine (Cortes and Vapnik, 1995), K-Means clustering (Bock, 2008) and Convolutional Neural Networks (LeCun *et al.*, 1998). Specialised areas were also covered such as Deep Learning for human activity recognition and Machine Learning for computer vision. One of the lectures consisted of a mathematics primer for Machine Learning covering probability, differential calculus and vector and matrix algebra. When describing these mathematics concepts the module leader frequently described them as “*simple*” or “*straightforward*.”

Students freely asked questions throughout the lectures for example one student asked, “*what is a loss/hypothesis function?*”. On a number of occasions the module leader did not directly answer the question and instead advised that this will be explained later. The lecturer regularly checked with the students whether they had any questions and asked whether they were ok with the material that had just been covered. There was usually minimum response to this questioning, however students often waited to speak to the module leader individually either during a break or at the end of the lecture. The module leader also indicated at the

beginning of one of the lectures that they had received a number of questions via email on the Support Vector Machine (Cortes and Vapnik, 1995) and proceeded to recap this topic. Alongside the interaction between the students and the lecturer, the students often interacted with each other during the lectures and particularly during the breaks when they would discuss topics they were unsure about.

To begin the lectures, the module leader usually introduced the topics to be covered in that session as well as referring back to content covered in previous lectures and how these concepts were related. Throughout the lectures the module leader would also direct students back to questions which were asked previously. At the end of that day's lecture, the lecturer would summarise the topics which had been covered and provide suggested reading as well as introduce the topic for the next lecture. As this module was block taught, at the end of each week there was a review of all of the content covered that week and a chance for the students to ask any further questions they had.

Code examples and real-world examples were regularly interspersed within the lecture material. On a few occasions this was research that the lecturer themselves had undertaken which seemed to further engage the students. The lecturer also advised the students to return to the theory of specific algorithms once they had implemented their code to solidify their learning and knowledge, to ensure they comprehended both the theory and practical aspects. The module leader often used diagrams and visual representation to display the difference between algorithms, for example a diagram was used to display the difference between the Multi-Layer Perceptron (Brownlee, 2016) and the Recurrent Neural Network (Hochreiter and Urgen Schmidhuber, 1997). The students were often required to work through a few questions within the lectures which were then discussed. Engagement with the questions was high, however when students were pressed for their answers there was usually minimum response.

A number of the slides within the lectures were very mathematics heavy, when explaining some of the algorithms. This may lead to difficulties for the students if they do not understand the notation. Due to the wealth of topics covered within this module, some of these concepts were covered fairly quickly, particularly Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997). This approach could lead to issues concerning the assessment if the

students are required to implement these algorithms for their coursework and the pace has been too quick for them to understand properly.

A number of interesting points were discussed by students during the breaks in the lectures relating to difficulties they were encountering within the module. These issues mainly pertained to the wealth of material being covered and mathematical/theoretical concepts. One student stated that they *“don’t understand enough to ask questions”*, another said that they were struggling to follow. One student said that they had studied A-Level Mathematics and still only understood 70% of the material covered in the previous lecture. Another student said that it *“seems so much compared to other modules, having to learn like 19 models.”* However, one student said that they *“usually find modules boring – not this one.”*

4.5.1.3 One-Minute Paper

Out of the three taught weeks of the Machine Learning module, the one-minute paper was deployed in the first two weeks as these contained the majority of the taught material and the guest lectures in the third week were not relevant to the assessment. This technique can quickly become laborious which can lead to declining completion rates (Stead, 2005). The one-minute paper was identical for both weeks apart from the content list of topics taught that week (see Figure 8 (p.72) for the one-minute paper for week one, see Appendix E for week two). The questions asked of the students were:

Q1: Which topic(s) did you find the most difficult this week?

Q2: Which topics are you still unsure on?

Q3: Any specific questions?

Q4: How confident do you feel in what you have learnt this week?

The final question, Q4, required the students to rate their confidence on a scale of 1 to 10, with 10 indicating they were very confident. Overall, 31 participants completed the one-minute paper, 16 participated in the first week and 15 respondents within the second week. The results from the one-minute paper can be categorised as pertaining to *troublesome topics* (from Q1 and Q2), *student questions* (from Q3) and *confidence* (from Q4).

Troublesome Topics

As shown in Figure 11, there were a number of topics which the students identified as being difficult to learn. The two most frequent subjects were the Support Vector Machine (Cortes and Vapnik, 1995) and the Multi-Layer Perceptron (Brownlee, 2016), followed by the Recurrent Neural Network (Hochreiter and Urgen Schmidhuber, 1997). Three out of the five topics identified as most difficult to learn were specific models, the further 2 were specific domain applications including Deep Learning for human activity recognition and Machine Learning for computer vision.

Comparing the results from Figure 11 and Figure 12 we can see that all of the subjects listed as difficult were also most frequently the topics which the students were still unsure of. The Convolutional Neural Network (LeCun *et al.*, 1998) was raised by some participants for Q1 as a subject they had difficulty with, however this was not raised as often as the other algorithms in Figure 11. However, Convolutional Neural Networks (LeCun *et al.*, 1998), alongside Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997) were the most frequent responses for Q2, with 22% making up nearly half of the overall responses. The frequency response relating to the Convolutional Neural Network (LeCun *et al.*, 1998) indicates that although a number of students might not have found that specific material covered in the lecture difficult, they still feel like they require further tuition in this area to fully grasp the architecture.

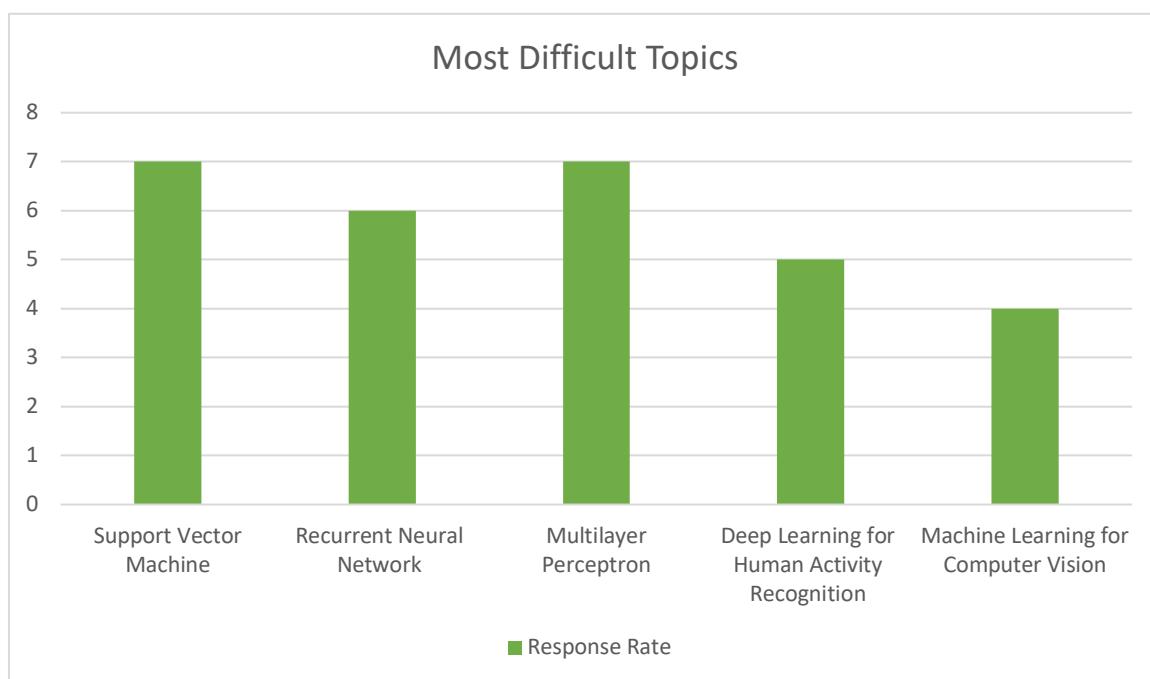


Figure 11: Most frequent responses to Q1: Which topic(s) did you find the most difficult this week?

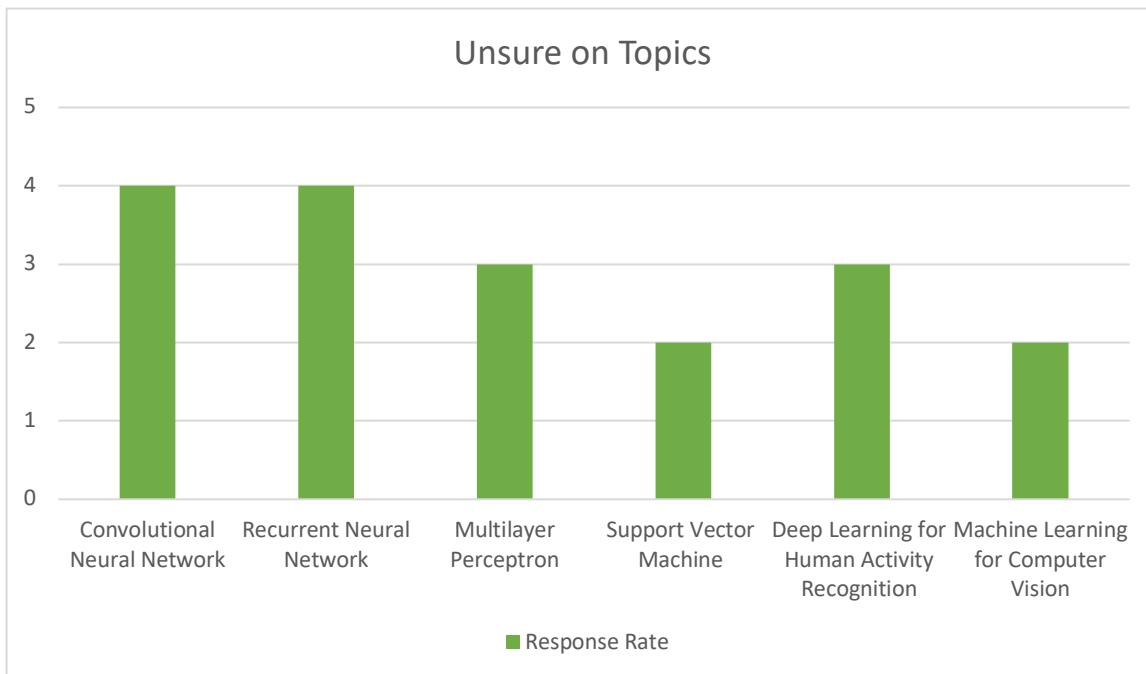


Figure 12: Q2: Which topics are you still unsure on?

Student Questions

Many students took the opportunity provided by the one-minute paper to ask specific questions. To understand the nature of the questions being asked, content analysis was undertaken. Two common themes emerged, students were unsure on the backpropagation algorithm (Rumelhart, Hinton and Williams, 1986) and they did not fully understand feature engineering. 57% of respondents mentioned these two topics. Feature engineering is a key process within any AI project. It is the “task of improving predictive modelling performance on a dataset by transforming its feature space” (Nargesian *et al.*, 2017). The backpropagation algorithm (Rumelhart, Hinton and Williams, 1986) is a vital component of many Deep Learning architectures and is important to understand not only the workings of many models but to also successfully implement these algorithms.

Other questions asked pertained to dropout, a technique for preventing overfitting in Deep Neural Networks (Srivastava *et al.*, 2014), the Gaussian distribution and the advantages and disadvantages of different models including the questions: “*is the multiclass logistic regression like discriminant analysis?*” and “*are there suitable applications for the different models – pros/cons*”.

Confidence

The average confidence level from week one from the one-minute paper responses was 5.1 out of 10. By week two the average confidence level rose to 6.8 out of 10. These results show that the overall confidence level increased between the two weeks. However, the confidence level was not exceptionally high. As the participants were asked to rate their confidence in what they had learnt that specific week, the rise in confidence could be a result of the material being covered in the second week being easier. Or it could be a consequence of the students becoming more confident overall within this domain.

4.5.1.4 Post-Module Questionnaire

The response rate for the post-module questionnaire (in Appendix F) was relatively low compared to the pre-module questionnaires at 15%, with only 7 students participating. One of the possible explanations for this low response rate was that the questionnaire was online, and the students were made aware of this around the same time as their assessment deadline.

The participants were asked to rate which topics, from the list provided, they found the most challenging, on a scale of 1 to 6 (6 = most challenging). The topic list was provided to the students, these included the key topics from the module catalogue. The topics which participants found the most difficult included backpropagation, 50% of respondents placed this at a 5 on the scale. Deep Learning was also identified as challenging with 66% of respondents placing this at 4 or over on the scale. Participants were split on mathematics, with 33% rating this as a 1 on the scale and 33% rating it as a 4. The same result was reported for dimensionality reduction. The participants were also asked which aspect of the module they found most difficult relating to either the theory or practical elements. 66% of participants responded that they found the practical aspects more difficult. Within the practical sessions the students were given a number of tasks pertaining to the learning algorithm covered in the lecture which they had to complete using Python.

The participants were asked which resources they used to aid with their learning, including textbooks, websites and online courses. The scale for this response was 1 to 4, 1 = never, 2 = rarely, 3 = sometimes, 4 = often. The resource which had the most frequent and highest score on the scale was websites, with 50% of respondents placing this at a 3 on the scale. The participants were also asked to rate how useful they found these resources on a scale of 1 to

5, with 5 being the most useful. Websites were the highest rated with 66% of respondents rating this as 4 or above, followed by online courses, with 33% of respondents placing this at a 5 on the scale.

Participants were asked how useful the differing learning strategies they used were. They identified note taking and practical exercises as the most useful. They determined that other techniques such as goal setting and reflection were not useful for them. The respondents were also asked to rate how confident they felt in applying Machine Learning upon module completion. This question used a scale of 1 to 5, with 5 being extremely confident. The results were mixed for this question, with responses at all points of the scale apart from 1 which indicated very low confidence. However, the most frequent response by a small margin was 5 with 33%. The intermediate points on the scale all had a frequency response of 17%.

To complete the questionnaire, there was a section where students could provide any additional comments. Half of the respondents took this opportunity to provide further information. Two of these comments pertained to the mathematical/theoretical aspects of this course. One of these comments stated that they had a strong statistics background so did not have difficulty with the amount of mathematical notation included in the slides, however they said they knew some students on the module found this challenging. Another respondent said that they found that *“the underpinning mathematical/statistical theory was left unexplained or not touched on enough.”* They further stated that as a consequence of this *“...the deeper mathematical theory underpinning these developments e.g. convolutional layers, I didn’t really fully understand.”* Another comment from a respondent challenged the amount of material covered in the module saying that *“the three-week module felt too compact to fully digest the information and this made it hard to keep up.”*

4.5.2 Artificial Intelligence Module

The Artificial Intelligence module at University B is taught at undergraduate level over a period of twelve weeks through a mixture of lectures and practical sessions. University B is a former polytechnic, considered part of the “post 1992” universities, which transformed UK higher education through widening participation (Scott, 2012). Content on this module includes knowledge representation and reasoning, search, and natural language processing. The second half of the module has a focus on Machine Learning, covering topics such as clustering,

neural networks and the application of Machine Learning techniques for computer vision. Assessment for this module consists of a portfolio based around two differing prototype scenarios; pathfinding with planning and search or Machine Learning to solve a real-world problem (Appendix H). Once the students have chosen their scenario, they are required to produce a report outlining the problem they are attempting to solve and then their coded solution. The students are also required to submit their code from the practical sessions.

4.5.2.1 Observation

Observation was completed with one cohort of students, with 6 sessions being observed. Students participating in this module were third year Computer Science undergraduate students and this module was taught within the first semester in a linear fashion. The same observation guide (Appendix D) was used for all of the institutions participating in the case studies, therefore the categories from the content analysis are the same. Average attendance for the lectures was around 25 students and lectures lasted one hour.

The topics covered within the 6 lectures observed are shown in Table 15. When introducing each topic in the lecture, the scope of this concept was defined as well as the historical context and example application areas. For the session on artificial neural networks, students were first introduced to how humans learn so they could build a mental model and process the new concepts presented to them in this lecture. The students appeared engaged and enthusiastic about learning this specific topic.

Lecture Topic	Content
Uncertain Knowledge and Reasoning	Probabilistic networks, Bayesian networks, game theory, Nash equilibrium
Introduction to Machine Learning	Reinforcement learning, classification, regression, supervised and unsupervised learning, building Machine Learning systems
Clustering	Hierarchical clustering, K-Means, nearest neighbour, outliers
Artificial Neural Networks	How the brain works, Perceptron, backpropagation
Introduction to Natural Language Processing	Finite State Machine, N-grams, Hidden Markov model, probabilistic grammars
Computer Vision	Image processing, principle approaches, example application areas

Table 15: Topics from the observed sessions of the Artificial Intelligence module

The amount of student and lecturer interaction was high within these sessions through the inclusion of group activities. This enabled active learning and appeared to increase student engagement as the students knew they were to be questioned on the content of the lecture. The inclusion of a group task also appeared to further increase interaction as the students were more forthcoming with questions on concepts which they did not understand. A number of student questions related to the mathematical notation on the presentation slides, one student was unsure about the notation for Bayes theorem.

Within the lectures, the lecturer regularly discussed their own personal experience and research within this domain as well as providing analogies to concepts familiar to the students. Real-world examples and applications of these technologies were also discussed. In discussion with the lecturer after the session, the lecturer advised that they try to link the practical exercises to the task within the lecture to incentivise the students to complete the practical exercises as attendance for the practical sessions was usually lower than the lectures. One of the points raised by the lecturer in the discussions after the lectures was the wealth of information and potential content to fit into twelve weeks and the importance that the students get to create something purposeful out of this knowledge.

4.5.2.2 Post-Module Questionnaire

The post-module questionnaire used for the Artificial Intelligence module at University B was identical to the one used for University A apart from the course content listed (Appendix F). Students were provided with the link for the online questionnaire towards the end of the module. Sixteen students completed the questionnaire, this was a response rate of around 64%. Relating to the demographics of the participants, all of the questionnaire respondents were male, this correlated with the data from the observations in that there were no female students on the module. The majority of the students (at 75%) were under twenty-eight years of age.

As with the questionnaire for University A, the participants were asked their highest mathematics attainment level. The majority of the respondents said their mathematics attainment level was GCSE or equivalent (at 50%), the next highest was A-Level at 31%. None of the respondents rated their confidence in their mathematics ability below a 5 on the scale of 1 to 10 (10 was exceptionally confident). The mean confidence level in mathematics skills

was 7.5 out of 10. The majority of the participants (at 68%), identified their programming skill level as novice. The mean confidence level relating to the student's confidence in their programming skills was 6.13 out of 10.

Within the questionnaire, the students were asked to rate which topics, based upon the module catalogue, they found the most challenging on a scale of one to five (five was equivalent to the most challenging). The most frequently highly scored topics on the scale included search algorithms, K-Means algorithm (Bock, 2008), Artificial Neural Networks (Abiodun *et al.*, 2018), Multi-Layer Perceptron (Brownlee, 2016) and backpropagation (Rumelhart, Hinton and Williams, 1986). Overall, 56.3% of the respondents found the practical aspects of the module the most difficult and 43.8% found the theoretical parts of the module more difficult.

The resource that participants listed they used the most and found the most useful were websites. The responses revealed that they either did not use or rarely used textbooks and online courses. From the learning strategies listed on the questionnaire, the most highly rated included note taking, practical exercises, goal setting and reflection. Participants were also asked how confident, on a scale of 1 to 5 (5 = very confident) they were on completion of the module of applying AI/Machine Learning techniques. The most frequent response on the scale was 3 with 50%, the next most popular response was 2 with 18.8%. The respondents were given space within the questionnaire to note any additional comments, but again, the majority of participants chose to leave this blank. However, two of the respondents provided positive feedback regarding the module, advising that they now understand AI. One respondent stated that more practical exercises would have been beneficial to them. Another participant advised that they had felt overwhelmed at first when they started the module.

4.5.3 Machine Learning and Computer Vision Module

The Machine Learning and Computer Vision module is offered at undergraduate level at University C which is also a “post-1992” university. The module is taught through a mixture of lectures and practical sessions, which alternate every week. The module content includes supervised and unsupervised learning as well as an introduction to the mathematical concepts underpinning this domain, such as probability theory. The module also has a focus on computer vision, covering application of Machine Learning techniques for example, to

biometric systems. There were no module pre-requisites, however students were expected to have some prior programming experience. Module assessment was purely coursework based.

4.5.3.1 Observation

Observation was undertaken over four sessions with one cohort of students. This module is a linear module and runs over one semester. Average attendance for the lectures was around 15 students, the majority of students were male and there appeared to be a wide age range of students. The lectures were split over two hourly sessions divided by a short break, however in the second session the lecture resumed within a different building. Over the course of the observation sessions, it appeared that the break and move to another building appeared to negatively affect student concentration and in more than one session, students who were present in the first half of the lecture did not return for the final part of the lecture.

To start each lecture there was a recap of the previous material covered and an outline of what was going to be taught in the session. The use of both real-world and code examples was used to assist with the explanation of concepts such as Bayesian classification (Theodoridis, 2020). Mathematical notation was included in some of the lecture slides to define some of the concepts, however, time was dedicated to explaining this notation and then questioning the students to ensure they understood. Overall, there was a lack of interaction between the lecturer and students as well as between the students themselves. The lecturer would often ask the students if they had any questions on the material being covered but this would rarely gain a response. The lack of interaction between peers was also notable, many of the students sat alone and did not seem fully engaged in the lecture. It perhaps would have been beneficial to include some form of exercise or group task to attempt to increase engagement and interaction.

4.5.4 Key Themes from Case Studies

There were a number of key themes under investigation within the case studies, these included student demographics and educational background, what level these modules are offered and how they are currently taught. There was also a focus on identifying the student experience when undertaking modules within this domain and determining topics and specific areas in which they are encountering difficulties.

Two of the modules within the case studies were offered at undergraduate level (Universities B and C) and one at postgraduate level (University A). All of the modules were taught through a mix of lecture and practical sessions. Unfortunately, specific student demographic information was only available for two of the participating modules, Machine Learning (University A) and Artificial Intelligence (University B) where students completed a questionnaire. For both of these modules the majority of students were under the age of twenty-eight and gender demographics were in line with other computing subjects in that there was a lack of female representation. Student educational background differed for the two modules, which was to be expected as one of the modules was offered as part of a masters course, therefore the students were required to have a previous degree within a scientific domain. The highest level of mathematics attainment also differed for these two sets of students. The students on the postgraduate module overall had higher mathematics attainment than the undergraduate group. However, the students with the lower mathematics attainment had a higher mean self-reported confidence level. Both sets of students rated themselves as novice programmers. The students completing the postgraduate module had the higher mean self-reported confidence level.

The students on the postgraduate module (University A) completed a pre-module questionnaire, where 71% of the 21 respondents expected that they would struggle with the theoretical aspects of the module. The students were then asked the same question on the post module questionnaire to determine which aspect of the module they actually found the most difficult. The majority of respondents actually found the practical side of the module more complex than the theory. The students on the undergraduate module who were completing the post module questionnaire (University B) were also asked which aspect they found most difficult, they also found the practical side of the module harder than the theory.

All of the modules covered supervised and unsupervised learning, including specific topics such as classification, regression, clustering and artificial neural networks. These modules also covered specific domain application areas, most notably computer vision was covered in all of the modules. The observation of the sessions allowed for commonalities in pedagogical approaches to be identified such as the importance of both real world and coding examples to enable the students to contextualise the theoretical underpinnings of the Machine Learning

models they are learning. The importance of student engagement and the use of activities was also noteworthy in that this fostered interaction not only between the lecturer and students leading to a more open atmosphere, where students were more forthcoming with questions or difficulties they were encountering but this also aided with peer to peer cooperation and communication. Through observation and analysis of all questionnaires, specific areas in which students were encountering difficulties were established. These included Deep Learning, specifically the backpropagation algorithm and artificial neural networks. Issues with the use of mathematical notation to explain specific Machine Learning modules also became apparent during the observations as the majority of student questions usually related to the requirement of an explanation of what this notation meant.

The questionnaires for the two modules which were participating in this form of data collection (Universities A and B) also contained questions relating to study and learning strategies the students were using. For both institutions' students raised websites as their most useful and frequent additional resource. Both sets of students used notetaking and practical exercises as a learning strategy. However, the undergraduate students (University B) were also using more holistic learning strategies such as goal setting and reflection.

4.6 Discussion

The key findings from all of the data collection methods outlined in this chapter are reviewed in the following section to determine the current range of education provision and student and lecturer experiences within this domain.

4.6.1 Current Provision

Within the universities reviewed as part of the online module review, all of the international institutions offered not only modules within this domain but full degree specialisms. This was not the case with the UK universities and suggests that there may be a potential for growth within the UK HE sector. Increasing literacy in AI throughout the population is a key recommendation for the UK government National AI Strategy as well as a suggestion that the demand for graduate level places on such courses requires a five to tenfold increase in admission numbers (UK AI Council, 2021). A review of the AI modules offered within the institutions reviewed as part of this study revealed an increase in provision through the

checkpoints of this data collection method. One potential avenue of inquiry is to repeat this study to determine whether this is a trend and universities are continuing to extend their AI course provision or whether this plateaus.

Out of the 25 universities reviewed as part of the online review of modules, 22 offered some form of Machine Learning module, this excludes Advanced Machine Learning modules, therefore these were considered introductory modules. Only 4 of the institutions offered an Advanced Machine Learning module. The majority of the Machine Learning modules were offered at undergraduate level. The most common content included classification, regression, clustering and neural networks. Specific algorithms included Naïve Bayes classifier (Gandhi, 2018), Support Vector Machine (Cortes and Vapnik, 1995) and Decision Trees (Quinlan, 1986). Content relating to Deep Learning was also prevalent within these courses as this is a subdomain of Machine Learning. However, some of the Machine Learning modules did not cover any form of Deep Learning. Deep Learning content included Convolutional Neural Networks (LeCun *et al.*, 1998), Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997) and the backpropagation algorithm (Rumelhart, Hinton and Williams, 1986). Only 10% of the Machine Learning modules included content relating to the ethics, legal and social issues associated with this technology. This is concerning and demonstrates a need for the wider adoption of ethics inclusion within these courses to ensure students become responsible practitioners.

Nine of the institutions offered a specific Deep Learning module. These were mainly offered at postgraduate level. The most frequent topics for inclusion in this module included Convolutional and Recurrent Neural Networks, Generative Adversarial Networks (Goodfellow *et al.*, 2020), regularisation, optimization and the backpropagation algorithm. Fifteen of the universities offered a specific Artificial Intelligence module and these were mainly offered to undergraduate students. 68% of these modules covered Machine Learning, including some form of Deep Learning.

The results from the lecturer questionnaires indicated the most commonly taught topics included linear and logistic regression (Worster, Fan and Ismaila, 2007), K-Means (Bock, 2008), Principal Component Analysis (Jolliffe and Cadima, 2016), Support Vector Machine (Cortes and Vapnik, 1995) and Bayesian Machine Learning techniques (Algorithmia, 2020). The

participants also identified the topics which they thought were pivotal to teach within these modules and the most frequent response was Deep Learning. This was reflected in the module content as all participants taught some form of Deep Learning.

All of the modules participating in the case studies taught some form of specialism including computer vision and natural language processing. Two of the modules in the case study also provided a mathematics primer or introduction to the mathematical concepts of this domain which included instruction on probability, differential calculus and vector matrix algebra. This content was covered within a single session which is fairly limited considering the importance of understanding these concepts, particularly within the Machine Learning module (University A) which was mathematics-intensive.

Out of the modules reviewed, Machine Learning and Artificial Intelligence modules were mainly offered to undergraduate students. Deep Learning and Data Science modules were primarily offered at postgraduate level. There was also variation relating to the content offered across modules within this domain, where some modules covered principles of AI, Machine Learning and Deep Learning rather than focussing on a specific area such as just Deep Learning. The inconsistencies in module content indicate a need to define the pedagogical content knowledge relating to AI/Machine Learning which would advance the “knowledge base of specialized teaching knowledge” (Hubbard, 2018). The QAA (2019) outline that the field of Computing incorporates both Data Science and AI which includes probabilistic Machine Learning techniques. However, the subject benchmark statement does not specifically list any core algorithms to be taught.

4.6.2 Lecturer and Student View

The use of questionnaires and interviews enabled comparison between the views of lecturers and students participating in a module within this domain and facilitated understanding of where both sets of views align and where they differ. The main area under investigation was understanding the issues faced by both sets of participants based on their experiences with Machine Learning.

The lecturers participating in this round of data collection identified that students have difficulty with the mathematical aspects of their module, and this incorporates into a wider

issue pertaining to the theoretical aspects of this domain. Other issues with the cognition of the domain theory include difficulties understanding terminology and matching an appropriate algorithm to a dataset. Students on the Machine Learning module at Newcastle University (University A) participating in the case study initially aligned with the lecturer view in that 71% expressed that they expected to find the theory harder than the practical aspects. However, within the post-module questionnaire this cohort were asked which aspect they actually struggled with the most and 66% said the practical aspects. This result also correlated with the findings from the Artificial Intelligence (University B) post-module questionnaire where 56.3% of students found the practical aspects harder. Although this cohort were much more split than the Machine Learning module as 43.8% had greater difficulty with the theory.

Topics which students identified as difficult to learn included Deep Learning and specific domain applications including human activity recognition and Machine Learning for computer vision. Specific algorithms included the Support Vector Machine (Cortes and Vapnik, 1995), Multi-Layer Perceptron (Brownlee, 2016) and Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997). The participants from the Machine Learning module were split on whether they found the mathematics difficult or not, indicating that some students had issues with the mathematical elements.

An issue that was raised by both lecturers and students was the amount of material that was covered within these modules. One lecturer advised that they struggle to fit the wealth of content into twelve weeks, another lecturer informed students at the start of the module that their module is “very challenging.” This was confirmed by students within the case studies as they found that a lot of material was covered. This particular module from university A covered Machine Learning, Deep Learning and specific application areas. Within the questionnaire, lecturers also disclosed feedback that they had received from students saying that the pace was too fast. They also advised that some students could become overwhelmed within lectures. This finding suggests that content to be covered on these modules needs to be reviewed and potentially limited to ensure that the cognitive load is appropriate.

4.6.3 Educational Background

Within the scope of this round of data collection, areas relating to educational background covered prior mathematics attainment, self-determined programming skill level and

perceived confidence in these skills. Student demographics and course pre-requisites were also included within this category.

Information pertaining to gender demographics, captured through the case studies, was in line with Computing in the UK in general in that there was and continues to be a lack of female representation (HESA, 2019, 2021). The average age for students on these courses was under the age of twenty-eight. There was some commonality relating to module pre-requisites in that there was usually some form of mathematics requirement. Linear algebra, calculus, statistics and probability were the most common specifications when there was some form of mathematics requirement. Prior programming experience was also an explicitly stated requirement for five modules. The prevalence of these pre-requisite specifications indicates that these modules are programming and mathematics intensive. In many of these cases, the prior knowledge required in both programming and mathematics were not formal qualifications or pre-requisites specified in module descriptions, therefore there may be variation in knowledge level within these cohorts of students, leading to some students having difficulties with these particular aspects. It also indicates that some module leaders may be assuming that students already have these skills and are therefore unaware of the differing educational levels. Students who do not have this prior knowledge may need to carry out additional learning to be able to cope with the content.

Within the case studies, highest mathematics attainment level differed within the Machine Learning and Artificial Intelligence modules. The majority of students on the Machine Learning module had up to A-Level Mathematics, whereas students on the Artificial Intelligence module had up to GCSE Mathematics. However, the students on the AI module had a higher mean confidence level in their mathematics ability. Student's metacognitive ability can be negatively affected by overconfidence and anxiety (Erickson and Heit, 2015). The study by Erickson and Heit (2015) aligned with the unskilled and unaware view of metacognitive judgement in that students often exaggerated their over-confidence in mathematics comparable to their performance. They found that this over-confidence also extended to "postdictions" for how well students thought they had performed. Students on the AI module reported that they found the practical aspects of the module harder than the theory, which is intrinsically linked with mathematics, but this was by a much smaller margin than the students on the Machine Learning module. The students on the AI module also rated their overall confidence in their

post-module ability to apply what they had learnt lower than the Machine Learning students. These findings may indicate a need for strategies to improve student metacognition to enable students to better self-monitor and self-regulate to gain more precise cognition of their ability. Erickson and Heit (2015, p.8) suggest self-testing to help improve self-monitoring, especially within a less stressful environment which will reduce distraction from working memory stores, enabling students to “perform closer to their actual knowledge level” and “highlight gaps in their knowledge” to enable the students to formulate better study strategies.

Both cohorts of students participating in questionnaires identified mainly as novice programmers. These students also had very similar self-reported confidence levels at around 6 out of 10. Both sets of students also identified the practical aspects of the module as being more difficult than the theory. One potential reason for this may be an issue highlighted by one of the lecturers participating in this study, that students struggle to identify the right model to use for specific datasets. Providing students with additional material relating to the appropriate applications of specific algorithms may help.

4.6.4 Learning and Teaching Strategies

The teaching strategies identified through this round of data collection established that AI/Machine Learning modules follow the general Computer Science convention of teaching this subject through a mixture of lectures and practical sessions. Pedagogical content knowledge in Computer Science indicates that this is inherently an active learning discipline and that students develop their problem solving skills through exercises and activities (Koppelman, 2008). This correlates with the view from one of the lecturers who advised that they think students learn better by doing something. Active learning through the use of group activities and questioning/problem solving during the lectures was prevalent throughout the observation sessions. As well as increasing engagement with the lecture material, the use of this teaching strategy also appeared to strengthen communication with the lecturer, appearing to make students more forthcoming with questions as well as improve peer to peer interaction. The level of interaction, or lack of, was noticeable in Universities A and C, where group activities were not frequently used, although individual problem solving was. This may have contributed to the reluctance of students to both answers questions posed by the lecturer as well as asking questions in front of peers due to a lack of confidence. This lack of

confidence may also have resulted from comments made by the lecturer in University A that some of the content was easy.

The use of real-life examples and code demonstrations were also common teaching strategies employed. One lecturer participating in the case study advised their students to revisit the theory after implementing the code examples to solidify their complete understanding of the model. Reiterating and recapping prominent models was also commonly used as well as discussing how particular models are disparate or related to one another. One lecturer within the case study would outline the full scope of an algorithm, including the historical context and example application areas before providing specific implementation details of how the algorithm works. For example, when the students were learning about Artificial Neural Networks (Abiodun *et al.*, 2018), they were first introduced to how humans learn and how this influenced ANNs, helping the students to build a mental model through a concept they could relate to.

One potential issue identified in two of the institutions participating in the case studies was the use of mathematics heavy slides to explain specific methods and how these work. Introducing students to the mathematical theory for these models is essential to provide a basis of how these algorithms work, however this could become problematic if it is the only method of instruction as some students may not have the requisite mathematics knowledge or may suffer from mathematics anxiety which could limit their comprehension.

The most common learning strategies identified through the case study included note taking and practical exercises, however the undergraduate cohort of students from University B were also using approaches such as goal setting and reflection which have been shown to help prepare them for self-regulatory practices (Nilson, 2013a). The most commonly identified resource to assist with learning, from both lecturers and students, was online resources. There is a wealth of online resources pertaining to AI and Machine Learning, however it can sometimes be overwhelming for students and staff to identify reliable sources or find information which is at the relevant educational level.

4.7 Summary

This chapter has outlined the statistical and qualitative results from the data collection methods outlined in the Chapter 3 (p.61) to provide answers to the research questions outlined in Chapter 1 (p.1). The online review of modules outlines the prevalence of AI/Machine Learning within the Computing HE sector, however there is still potential for growth, particularly within the UK. The majority of Machine Learning modules were offered at undergraduate level and often required pre-requisites pertaining to mathematics, particularly linear algebra, calculus, statistics and probability. The data from the lecturers teaching this domain identified that the students often struggle with the mathematics and theoretical aspects of these modules. The MSc cohort of students participating in the case study from university A also expected to struggle in particular with the theory. However, the questionnaires completed after the module indicated that the students actually found the practical side of this domain harder. This disconnect between the lecturer and student view may indicate an issue with student metacognition in that they are not accurately identifying their theoretical understanding. Or it may indicate that only some students are struggling with the theory/mathematical elements.

One potential method to help improve student metacognition, particularly related to their mathematics knowledge is the use of self-testing as this can improve self-monitoring skills (Erickson and Heit, 2015). Incorporating self-testing into a distance learning online tool has also been shown to be a promising strategy to alleviate mathematics anxiety (Iossi, 2007). Coupled with the fact that online resources were identified by both lecturers and students as the most common resource for additional learning, an online learning tool which incorporates specific mathematics skills for Machine Learning seems a potential method to boost students' mathematics skills.

The finding that both students and lecturers were having difficulty with the wealth of content to cover and learn within these modules, often covering AI, Machine Learning and Deep Learning, highlights the need for a framework for an introduction to this domain to set out the essential knowledge. An online resource which encompasses the key areas, particularly of Machine Learning and Deep Learning, may be useful for lecturers to be able to point students to a resource which will provide students with the key knowledge. As well as provide additional tuition on the specific algorithms identified as troublesome, including the Support

Vector Machine (Cortes and Vapnik, 1995), Multi-Layer Perceptron (Brownlee, 2016) and Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997).

Chapter 5. An Online Learning Tool for an Introduction to Machine Learning

5.1 Introduction

Increasingly within Computer Science education, students are less frequently turning to textbooks to aid with their study and are instead using the internet to gain further information and guidance (Boese, 2016). However, the wealth of information online related to Machine Learning can be difficult to navigate and assimilate into the context of the student's own personalised learning. The importance and availability of online learning provision has been highlighted by the worldwide COVID-19 pandemic (World Health Organization, 2021). Higher education institutions which already offered some form of blended learning were at an advantage as they already had experience in hosting online material, whereas other universities had a lot of work to do in transitioning to an online environment (Crawford *et al.*, 2020). This transition to online learning posed a number of issues including "problems of delivery, staff expertise, and student engagement," as well as wider issues relating to "inequalities of access and outcomes in the new pedagogic spaces" (Peters *et al.*, 2020). Many practitioners within the field of higher education believe the effects of the pandemic will alter the state and provision of education within this sector permanently with the changes being linked to its digital transformation (Watermeyer *et al.*, 2020).

This chapter explores the design and implementation of a learning tool, called MetaLearning (links to tool in Appendix I) as outlined in RO2.b in Chapter 1 (p.6), which aims to create an environment in which the student can learn the most pertinent aspects of Machine Learning to gain an overview and comprehensive introduction to the field. This chapter details the specific aims of MetaLearning and the intended audience, the environment chosen to host the tool, the content and how this was chosen and presented. Finally, a use case to detail how the system can be used is discussed.

5.2 Rationale for Creating an Online Tool

Choosing to create an online tool to assist students in their learning of Machine Learning was motivated by a number of findings. The use of the internet is ubiquitous in the majority of courses as students often have to use an LMS [Learning Management System] to access course materials and upload assessments. Students are also likely to use online resources as part of their learning and the lecturer questionnaires showed that the course instructor will often

provide a list of online sources to assist the students in a particular topic (Chapter 4, Section 4.4.1, p.96). The benefits offered by this form of learning instruction were highlighted in Chapter 2 (Sections 2.8.1 and 2.9), particularly in alleviating mathematics anxiety and the potential of integration of the outlined mitigation strategies identified in the research objectives. The use of online learning was also identified in Chapter 2 (Section 2.9, p.55) as a method to improve student self-monitoring through the use of strategies such as quizzing with immediate feedback.

Within Computer Science students often seek advice from online discussion boards such as Stack Overflow (Stack Exchange Inc, 2022) and often find solutions for their assignments widely available (Boese, 2016). This form of plagiarism is a widely acknowledged issue in Computing courses (Sheard and Dick, 2011). Although it is not one of the aims of MetaLearning to try to prevent plagiarism, a self-contained and focused tool may prevent students from looking outside of the tool to other less reliable sources. The inclusion of strategies to improve student self-efficacy and confidence may also prevent students from feeling the need to seek solutions on the internet.

The rise of MOOCs, as discussed in Chapter 2 (Section 2.9.1, p.56) has brought about a wide variety of openly available courses on topics such as Artificial Intelligence, Machine Learning, Deep Learning and specific application domains. Both Udacity (2021) and Coursera (2021), two of the biggest providers of MOOCs, offer a range of courses related to AI. Udacity has a dedicated School of Artificial Intelligence (Udacity, 2020) offering courses such as AI Programming with Python, Deep Reinforcement Learning and AI for Healthcare. In 2019 the Introduction to Tensorflow for Deep Learning course had the highest enrolment till the mid-year when it was overtaken by the Introduction to Machine Learning course, which had over 125,000 people globally enrolled (Perrault *et al.*, 2019).

Coursera have over 845 search results related to Artificial Intelligence (as of October 2020) which have a range of difficulty from beginners' courses such as AI For Everyone to advanced courses such as Advanced Machine Learning with TensorFlow on Google Cloud Platform. The courses are offered by a variety of providers, from academic institutions, such as Stanford University, to prominent industry names such as IBM and Google.

A study conducted by Eckerdal *et al.* (2014) into computing academics opinions on MOOCs discovered what individuals in this field deem the main advantages and disadvantages of these types of courses. Identified advantages include affordability, limited restrictions relating to time and location of study and the ability to learn from eminent experts within the particular domain. Disadvantages were related to lack of consideration of pedagogical aspects such as design and assessment and a deficit in personalised interaction with the instructor. Opposed to traditional face-to-face courses, MOOCs also have a high dropout rate, with many participants failing to complete their course (Clow, 2013).

One of the main concerns discovered in the Eckerdal *et al.* (2014) study was the deficit of pedagogical consideration in the implementation of online learning tools. The Machine Learning tool to be created as part of this study will be constructed around the theory that learning is the main priority and that the technology employed needs to be the best option for the pedagogical aim of the learning (Curry, 2018). As identified in Chapter 2 (Section 2.6, p.37) students enrolled on Computing courses often enter the programme with varied educational backgrounds which can result in disparities in prior knowledge. The particular set of pre-requisites identified in Chapter 4 for courses within the AI domain require quite a high level of mathematics knowledge. Therefore, one of the aims of MetaLearning was to level up the cohort disparity in this knowledge domain.

Addressing the issue of mathematics anxiety and difference in students mathematical background was a priority due to the prevalence of mathematical concepts and level of knowledge needed to understand some of the ideas within AI. Another pedagogical aim and purpose of the online tool was to include content specific to the skills identified as lacking in graduates as identified in Chapter 2 (Section 2.6.2, p.40) including Machine Learning, data processing and ethics. Alongside the specific domain content knowledge there is a focus on addressing some of the barriers and difficulties students may encounter when learning AI as outlined in the RO1 (p.6). These barriers were identified as mathematics anxiety (Chapter 2, Section 2.8.1 and Chapter 4, Section 4.4.2) self-efficacy and metacognition with the aim of developing the students into greater self-regulated learners (Chapter 2, Section 2.8.4, p.54).

Due to the vast number of online courses available within the field of AI, it can be difficult for novices to navigate the wealth of information and discover which course is most appropriate

for their learning needs and current level of expertise. Many of the online courses are disparate in their offering for introductions to the AI discipline, for example there are separate introductory courses for Machine Learning, Mathematics for Machine Learning, Deep Learning and Artificial Intelligence. However, to get a good grasp of the overall AI domain, knowledge of each of these subdisciplines and subject areas is beneficial. Therefore, the proposed online tool will cover all of these topics to give a definitive introduction to this domain.

The content to be included within MetaLearning was determined by the findings in both Chapter 2 and Chapter 4. As previously discussed, students educational background, particularly in relation to mathematics and statistics can be a challenge for some learners, therefore a tutorial on mathematics and statistics knowledge specific to Machine Learning was to be included. Both the online review of modules (Chapter 4, Section 4.3, p.82) and the case studies (Chapter 4, Section 4.5, p.101) identified the most commonly taught concepts within Machine Learning, Deep Learning and AI modules as shown in Table 16. The concepts identified were often taught interchangeably within both specific Machine Learning and Deep Learning modules.

Commonly Taught Topics	
Machine Learning	Naïve Bayes Classifier, Support Vector Machine, Decision Tree, Linear and Logistic Regression, K-Means, Principal Component Analysis
Deep Learning	Convolutional Neural Network, Backpropagation, Recurrent Neural Network, Generative Adversarial Network,

Table 16: Commonly taught topics identified in Chapter 4

Alongside the algorithms, concepts such as regularization and optimization were also covered. The ubiquity of all of these topics meant that their inclusion in MetaLearning is essential to ensure that users gain a full introductory understanding of this domain. The findings from the case studies, particularly the one-minute paper (Section 4.5.1, p.111) identified topics which the students found troublesome, this included the Support Vector Machine (Cortes and Vapnik, 1995), Multi-Layer Perceptron (Brownlee, 2016), Recurrent Neural Network (Hochreiter and Urgen Schmidhuber, 1997) and Convolutional Neural Network (LeCun *et al.*, 1998). Therefore, content creation of these topics will have to consider the best methods to convey these complex topics. The interviews with the lecturers (Section 4.4.2, p.97) also

identified that learners had difficulty with the terminology, therefore a thorough explanation and clear indication of the terminology is important for inclusion. The finding from the online review of modules that only 10% of the modules included ethics within the content (Section 4.3.2, p.88) was concerning and highlighted the importance of inclusion of some form of instruction in this area within the online learning tool.

Providing an overview of the field of AI was deemed important to give users insight into what AI actually is and the many facets of this domain. This tutorial should provide context of where this domain originated and the preeminent research throughout its history. Content relating to the sub-disciplines of Machine Learning and Deep Learning will provide users with an introduction to concepts such as supervised and unsupervised learning, artificial neural networks and the Machine Learning workflow. With the inclusion of the AI tutorials, this should assist the learners in comprehending how the various domains sit within the AI umbrella.

Incorporating elements of AI and Deep Learning were important for inclusion in a tool which is mainly focused on an introduction to Machine Learning as there can be confusion, especially with newcomers to the field over what constitutes a Machine Learning model and what can be deemed Deep Learning. There is often a lot of overlap of terminology as well as content within modules/courses in this area, therefore MetaLearning aims to provide instruction on all of these topics to try and dispel this confusion. A reoccurring point from Chapter 4 (Section 4.6.2, p.124) raised by both students and lecturers was the wealth of content covered on these modules. Therefore, the inclusion of AI, Machine Learning and Deep Learning content may offer a solution to this issue as the lecturers could refer students to specific tutorials to complete outside of the lecture time to free up further time for pivotal topics such as the threshold concepts.

The intended audience for the tool is students studying at higher education level, mainly students in some form of Computing course, however the tool is also applicable to students in other domains who wish to learn more about Machine Learning. For example, the numerous applications of this technology may interest individuals studying Biology as there is a vast research area within medical diagnostics and the use of Machine Learning and Deep Learning to improve this area (Fujita, 2020). The complexity of the material is at a level which

is accessible for students who may not have a background in Computer Science or a strong mathematical background as was found in Chapter 4 (Section 4.6.3, p.124). Judging the difficulty level at which to pitch the material for the online tool was informed by the overall aim of it being an introductory course. However, opposed to the Machine Learning modules reviewed in Chapter 4 (Section 4.3.2, p.84) one of the key objectives of the online tool was that users were not required to have any prior knowledge, particularly within mathematics so that this online tool was inherently an introduction to all aspects of this domain.

Due to the pedagogical focus of the online learning tool and the importance placed on both the content and the mitigation strategies it was decided that the tutorials should be the key focus of this aspect of the research, therefore existing tools/software were reviewed for adaptation. To determine the most appropriate software in which to build the tutorials, requirements were drawn up relating to the functionality the adopted system should have in order to host MetaLearning.

5.3 Tutorial Requirements

The requirements for MetaLearning were elicited upon completion of the case studies, online analysis and consultation with AI module leaders detailed in Chapter 4. Alongside the domain content, inclusion of strategies to improve student self-efficacy, metacognition and self-regulation were identified for incorporation through Chapter 2 and 4, alongside techniques to alleviate mathematics anxiety. Specific techniques identified included retesting and self-paced learning potentially incorporated in some form of distance learning (Chapter 2, Section 2.8.1, p.50). It was also identified in Chapter 2 (Section 2.8.2, p.52) that the opportunity for students to have successful experiences relating to independent personal performance can reinforce individuals' sense of competence and consequently their self-efficacy.

There are many factors to consider when creating online tutorials to try and mitigate the high dropout rates prevalent within this learning domain. These include the significance of the content, the users ability with technology, however the most common factor is lack of consideration of design and usability (Zaharias and Poylymenakou, 2009). All of these issues were considered when planning and implementing MetaLearning. The following sections outline the design and technical requirements for the tutorials, this will assist in identifying the host technology.

5.3.1 Tutorial Design Requirements

The online courses offered, particularly through MOOC platforms offer a variety of different subjects at a number of levels from beginner to advanced. However, as discussed in the previous sections they often fail to properly account for different learning needs and do not meet the requirements needed to address the educational and cognitive requirements of newcomers to the field of AI. Therefore, the following design requirements were necessary for the MetaLearning tutorials:

R.1 Ability to host engaging multi-modality content

One of the most important elements of MetaLearning is ensuring that the tutorial content is engaging, pitched at an appropriate skill level and informative. Findings from Chapter 2 (Section 2.6.2, p.43) advised that it is best practice in Data Science education to include case studies and where applicable visualisation. The ability to incorporate various media mediums within the online tool is important to discover the best methods of conveying complex information including experimentation with visualisation and data sonification, which is the rendering of data and information into sound (Kaper, Weibel and Tipei, 1999). The use of these techniques will be trialled within MetaLearning tutorials to determine their effectiveness in conveying information.

R.2 Capability for Formative Assessment

One of the strategies to be employed within MetaLearning is a type of formative assessment through testing as a learning approach. The aim is to engage students with the content and to provide immediate feedback. Retesting and the opportunity for successful learning experiences were identified in Chapter 2 (Sections 2.8.1 and 2.8.2, p.47) as strategies to alleviate mathematics anxiety and improve confidence and self-efficacy. The use of low-stakes quizzing has been determined as a strategy to provide the users with a formative assessment experience.

R.3 Provide Learners with Feedback

Providing users with feedback is imperative to the use of testing as a learning tool as an approach to learning and improving student self-efficacy. The online tool should have the

provision for a number of differing options for presenting feedback, including acknowledgement of when an answer is correct or incorrect, this should then be expanded for specific questions to include an ‘advice’ section which provides a step-by-step guide to how the question can be solved. The ability to provide users with additional resources is also an important requirement within the feedback section so that the user can access additional provision if they are struggling.

Providing an overall mark for the tutorial, based upon the user’s success with each individual question will be an important aspect of the feedback for MetaLearning so that students are given an indication of their progress and to guide their further learning (Barbosa and Garcia, 2005). To assist students in their learning process to ultimately become self-regulated learners, initial suggestions will be provided to the students on how they can work to improve their marks through repetition of the tutorial or attempting other tutorials in the learning tool which may provide a foundational understanding.

5.3.2 Tutorial Technical Requirements

Ensuring that the tool is both accessible and relatively easy to use is of upmost importance to encourage users to actively engage and re-visit the tutorials. Providing information to the users on data use and storage is required to build trust, this is especially important within an e-learning system as the data gathered by the tool may be used to form assessment outcomes and grades. The technical requirements considered key for the adaptable existing software and tutorials include:

- Usability

Usability relates to a number of differing aspects of the tool, for example it is of high importance that the system is both intuitive and straightforward to use as the “usability of e-learning designs is directly related to their pedagogical value” (Zaharias and Poylymenakou, 2009). The unrestricted access to the tutorials and ease of finding the specific information was also of importance to ensure that any cognitive load, “the limited ability of the working memory to code information” (Kirschner and Hendrick, 2020, p.14) faced by users was through the learning materials and not through use and navigation of the system. This formed requirement:

R.4 User Interface must be attractive, consistent and responsive to the user

To guide the technical and design planning and implementation, the usability framework designed by Zaharias and Poylymenakou (2009) was used as shown in Figure 13. This framework identifies two strands of usability attributes for e-learning, general usability of the system and instructional design. General usability includes navigation, accessibility, and visual design. The instructional design aspect of the useability attributes includes interactivity and engagement, content and resources and media use. Both the general usability and the instructional design usability aspects have a focus on and can impact motivation as a learning dimension, this is particularly pertinent for online learners as the domain necessitates self-regulation and motivation to remain engaged (Artino and Ioannou, 2008). The motivation to remain engaged is a big focus for MetaLearning as the majority of students surveyed prior to undertaking a course within the AI domain had high motivation for their studies as outlined in Chapter 4 (Section 4.5.1, p.106). This formed requirement R.5:

R.5 Tutorials must be interactive and engage users to maintain motivation

The attributes for usability outlined by the Zaharisa and Poylymenakou framework (Figure 13) are consistent with Nielsen's usability heuristics (Nielsen, 1994) and offer a blueprint for design concerns. All of the points raised in the usability framework are to be considered during creation of the tutorials, for example to ensure that the visual design is consistent and that the tutorials are easy to navigate.

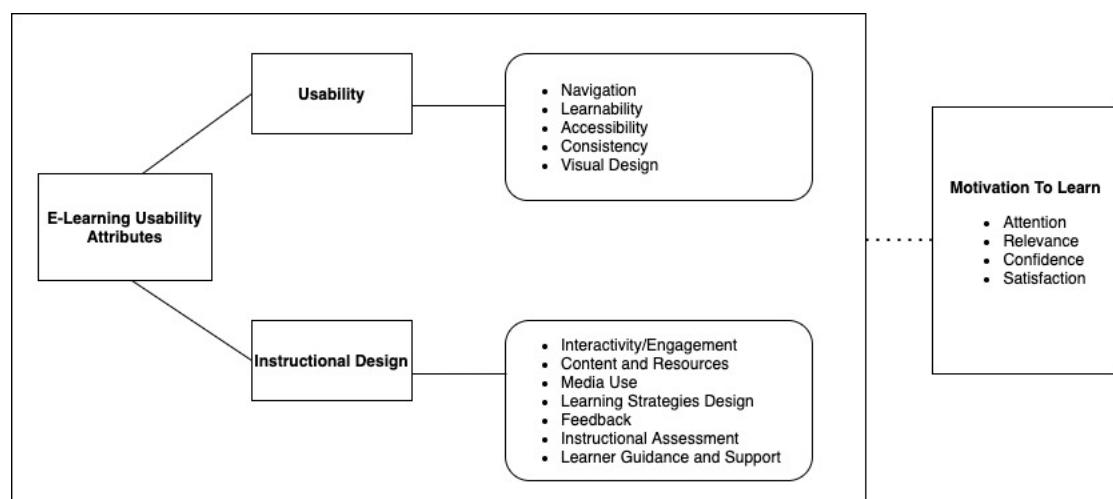


Figure 13: E-Learning Usability Framework (adapted from Zaharias and Poylymenakou, 2009)

- Deployment

Enabling users to easily access MetaLearning is imperative and consideration was given to how best to access the tutorials. Two potential methods for deployment have been considered, these include integration into a VLE (Virtual Learning Environment) or use as a standalone web page. This informed requirement R.6:

R.6 Ensure access to the tool is readily available

Both methods were deemed appropriate for differing cohorts of students based upon the case studies in Chapter 4 (Section 4.5). For example, students were regularly accessing resources through their VLE, therefore the tutorials should be made available using this method. It may also be beneficial to use a more targeted approach for particular cohorts by emailing them the links to MetaLearning.

- Privacy and Security

As outlined in Chapter 2 (Section 2.3, p.19) privacy and security of personal data is a big concern, this is especially pertinent in e-learning tools, where alongside personal data, information collected can extend to learning preferences and progress, the students background and behaviour as well as data relating to learning analytics (Ivanova, Grosseck and Holotescu, 2015). This informed requirement R.7:

R.7 No new personally identifiable data should be collected through the tool

Within MetaLearning, there will be no additional personal data collected apart from tutorial marks. For instances when the tutorials are incorporated within the VLE, the marks from the tutorials will be recorded and available to individuals with administrative access to that course. When the tutorials are used in standalone mode, the user will not be required to sign up or sign in with any personal details and their scores from the tutorials will not be recorded. This requirement also ties in with the use of distance learning to ease mathematics anxiety due to the anonymity of an online course.

5.4 Design

Upon completion of the requirements gathering process, design of MetaLearning could be considered. Based upon the requirements and the importance of the content for the online tool, existing systems were analysed for appropriateness for adaption and adoption for this suggested online tool so that focus could be centred on the domain content.

One of the other main design considerations was how the tool was to be used. To leverage all of the benefits afforded through online and distance learning, the tool is proposed for use as a standalone resource as well as for use in blended learning and flipped classrooms. This decision was informed by findings from Chapter 2 (Section 2.4.1, p.25) which identified that blended learning and in particular the use of technology has led to further democratization of learning as the flexibility enables students to fit their studies around other work and caring responsibilities. Findings from Chapter 2 (Section 2.6.2, p.44) also indicated that educators within this domain had a requirement for further teaching resources, specifically interactive tools which would enable them to incorporate active learning into their teaching. For use as a standalone tool, the user should be able to choose which tutorials they wish to engage with. There should be no set chronological order in which the tutorials have to be completed, users should be able to access whichever tutorial they identify as the most useful for them. However, reflecting on the variation in educational background of potential users, a number of tutorials should provide introductory overviews of the main domains, including AI, Machine Learning and Deep Learning, as well as tutorials which provide more in depth and complex content.

Blended learning, the “integration of traditional classroom methods with online activities” (López-Pérez, Pérez-López and Rodríguez-Ariza, 2011) was considered as an educational approach for MetaLearning as the tool can be used as a resource or additional activity. Additional activities can be used to reinforce the material taught within the lecture as well as improve motivation (López-Pérez, Pérez-López and Rodríguez-Ariza, 2011). The tool should also provide activities related to the process of learning such as knowledge surveys and the use of low stakes quizzing to help improve student competency and become more self-regulated learners.

MetaLearning should also be designed so it is applicable for use within flipped classrooms, for example the students would be required to complete a specific tutorial before the classroom session where any discussions could be facilitated based on any of the topics, or further instruction could take place on any specifics the students had difficulties with. The tutorials could also be used as a basis on which more advanced topics could then be built upon.

Creating MetaLearning requires consideration of not only domain specific content but also consideration of pedagogy and the identified barriers to mastery of the field of Machine Learning to build a learning tool with inbuilt mitigation strategies based on the findings from Chapters 2 and 4. Bringing all of these elements together required creation of a new introductory course for the AI domain as existing learning resources were not suitable for modification.

Based upon the requirements for the tutorials outlined in Section 5.3, varying existing educational systems were evaluated for applicability in hosting MetaLearning. Common VLE's were the main source of investigation including Canvas (Instructure, 2022) and Blackboard (BlackBoard Inc., 2022), however use of specific VLE's would limit accessibility and deployment as differing universities use varying learning environments. Therefore, other available e-learning resources were researched including Numbas (Newcastle University, 2015c), which was ultimately chosen as the host for MetaLearning due to it meeting the outlined requirements and the ease in which the mathematical aspects of the tutorials could be incorporated. Numbas is discussed in further detail in the following section.

5.5 Numbas

As the focus for MetaLearning was on the content and pedagogical aspects of the tutorials, deciding to use the readily available, easily accessible infrastructure of Numbas (Newcastle University, 2015c) allowed for targeted content creation. Numbas is an e-assessment system which is both free to use and open source, it was developed to replace commercial systems as a more modern alternative (Perfect, 2015). As an e-assessment tool the purpose of the system is to create tests consisting of a number of questions. Questions can consist of a number of elements including the statement, one or many parts and an advice section. There are a number of question types available including number entry, multiple choice and text entry. As a Numbas test is essentially a webpage, content can also include videos and

interactive graphics (Perfect, 2015). The system has previously been used in the HE sector for formative e-assessment and for early module diagnostic tests to identify student learning needs in regards to mathematics and statistics courses (Foster, Perfect and Youd, 2012). It has also been employed in schools to help students prepare for mathematics exams (Newcastle University, 2015b) and within commercial settings to help train new staff (Newcastle University, 2015a).

One of the benefits of the tool is the capability to provide users with timely, personalised feedback upon completion of both individual questions and exams. Feedback is incredibly important as it endeavours to minimise the disparity between “prior and current achievement and the success criteria” (Hattie, 2012) and has also been shown to increase motivation and self-esteem (Nicol and MacFarlane-Dick, 2006). It is also key within the consolidation phase of learning that the student becomes competent at receiving and reflecting on feedback (Hattie and Donoghue, 2016).

Another benefit of Numbas is the client-side design, this ensures that it is fast as there are no calls to external systems and it can also be easily distributed to users as either a standalone webpage or simply integrated into the university VLE (Perfect, 2015). However, there are drawbacks to Numbas running entirely on the client as this is not as secure as server-side designs. This may be an issue for high-stakes assessment as students who wish to, may find a way to cheat. However, as the proposed online tool is only for use as a learning and formative assessment resource, there is limited reason or repercussions for cheating or deception.

Integration of Numbas within a VLE allows exam scores to be recorded, ensuring that the author of the test can access the student’s results. However, this option is not available when running the test as a standalone webpage. This has many drawbacks including the lack of access to learning analytics to determine how students are using the tool and how much time they are spending on a tutorial etc. It also means that the author of the exam has no way of knowing which students have used the tool and what their results were from any tests they may have taken. However, these drawbacks are not an issue for the purposes of MetaLearning as outlined in requirement R.7.

Whether the student is using Numbas as a standalone webpage or within a VLE the process for use is exactly the same. When opening a Numbas exam, the user is presented with an overview of the number of questions within the assessment, the marks available and a 'Start' button which commences the test. Upon entering the test, the user is presented with a screen containing a side panel displaying the number of questions left, their score on each question and which questions are still to be answered. Within each question, the user is shown the question and provided with an answer box in which to input their answer. There are a number of options to choose from upon question completion including 'Submit answer' which provides immediate feedback on whether the user's answer is correct. The other options include 'Try another question like this one' which randomly selects another question within the same question category from a pool of similar questions and 'Reveal answers' which provides the user with advice on how they could have solved this question.

At any point the user can chose the option to 'End exam' which brings them to a performance summary page. Users can see a breakdown of their scores for each individual question as well as view a question review which displays the worked through solution for the question. Also displayed on this page is time the user spent on the exam and their total overall mark as a percentage. Throughout the development process and subsequent deployment phase of Numbas, user feedback has been implemented to ensure the interface for both students and exam authors is intuitive (Perfect, 2015).

As Numbas has already been deployed in a number of varying educational scenarios, outlined at the start of this section, Numbas already satisfies requirement R.4 as the system has already been part of a number of existing usability trials (Perfect, 2015).

5.5.1 How Numbas will be adapted

The suggested use of Numbas as a formative e-assessment tool will be integrated into the overall concept of MetaLearning. However, Numbas will be adapted to support individual tutorials to become more of an e-learning tool instead of a wholly assessment-focused resource. Keeping formative assessment as a feature of MetaLearning is important to address some of the issue's students may face when learning this topic including cognitive barriers such as low self-efficacy and confidence and attempting to mitigate these issues through strategies aimed at becoming more self-regulated learners. Part of the importance of

formative assessment is the opportunity for feedback and for students to comprehend and act upon the feedback they receive. This opportunity will be given to the students through the online tool as resubmission of answers and multiple opportunities to undertake the questions will be permitted. This will allow the student to “close the gap between current and desired performance” (Nicol and MacFarlane-Dick, 2006). Retesting was also identified in Chapter 2 (Section 2.8.1, p.50) as an approach which has been shown to reduce mathematics anxiety.

One of the main differences between the traditional use of Numbas and the use of it within MetaLearning is that the question setup will instead become tutorials which consist of a number of Numbas questions containing subject specific content. As shown in Figure 14, there will be an introductory tutorial page which will summarise the content covered as well as a summary of the number of questions within the tutorial and a pass percentage. This pass percentage is an indicator to the student of whether they require further tuition within this specific topic. At the end of the exam the user will be presented with guidance indicating that if they did not reach 50% then they are advised to re-do the tutorial and utilise some of the additional resources listed within the tutorial pages (as shown in Figure 17).

The user will have the option to change the display options (as shown in Figure 14) which satisfies requirement R.4, this functionality is already integrated into Numbas and allows the user to change the display and font size and display colours. Once the user chooses to start the exam they will be presented with a screen outlined in Figure 15. The individual tutorials will contain information pertinent to the specific topic covered and multi-modality material including visualisations. Additional resources will also be provided if the user wishes to investigate the subject matter further. As shown in Figure 15, the user navigates through the individual tutorials using the sidebar, this also displays to the user how far through the tutorial they are.

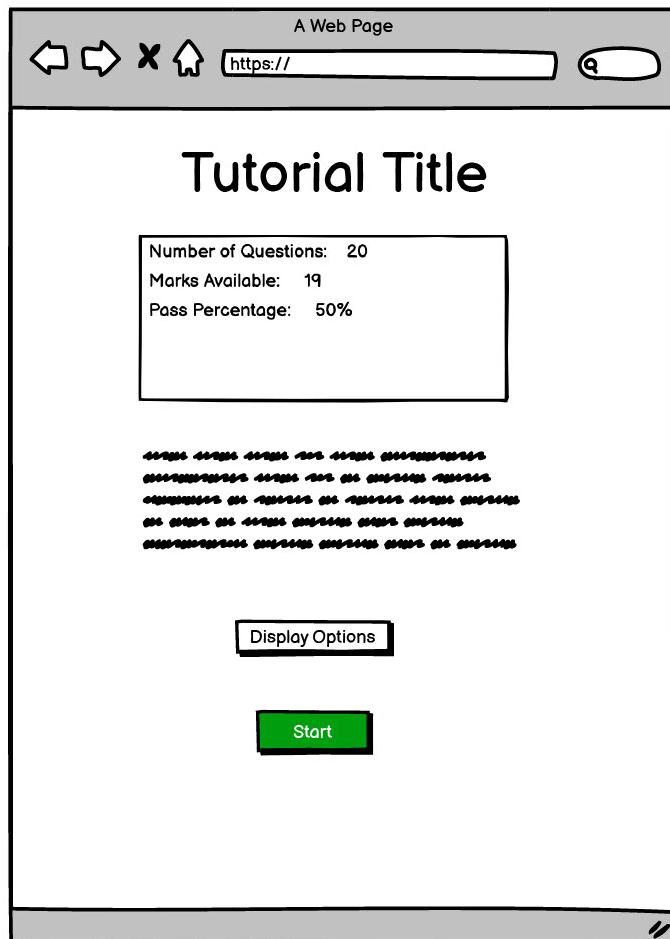


Figure 14: Introductory page for the tutorials

At the end of each content section, the user is posed with a question based on the content they have just read (as shown in Figure 15). The questions are low stakes and contribute to an overall module mark. The format of the answers required will vary to not only challenge the users recall through multiple-choice questions but also their understanding of the content they have just read, through free form text boxes and matrix entry boxes where the user will be required to undertake calculations to provide an answer to the question. Although the tutorials will still be labelled as an exam and questions, as inherent to the Numbas system, each 'exam' and 'question' will be a tutorial within a specific sub-domain of AI.

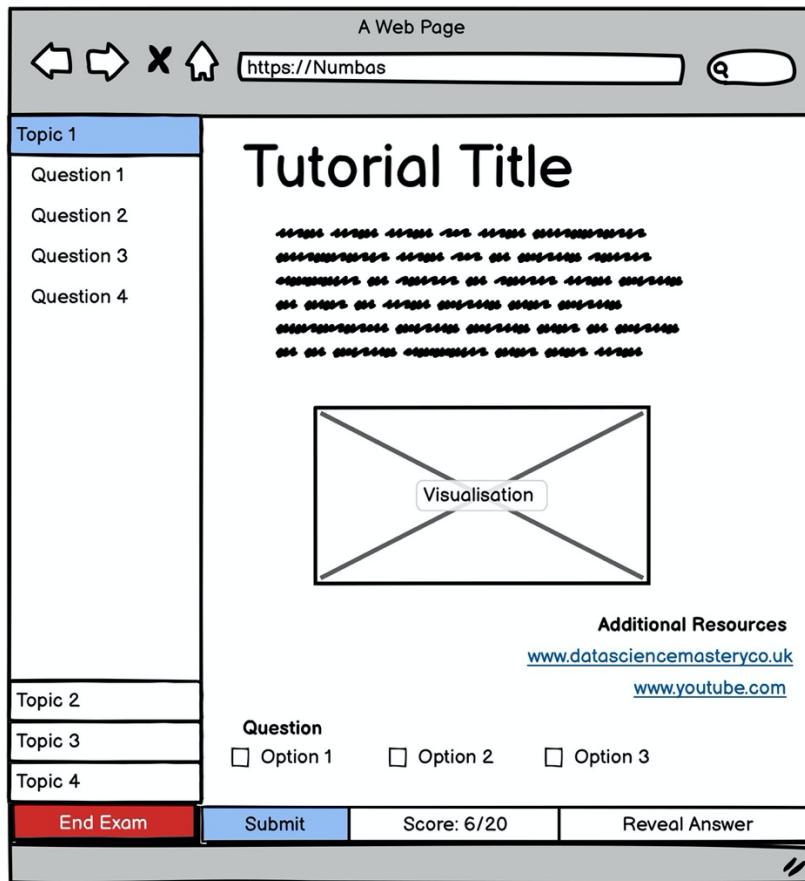


Figure 15: Design of the individual tutorial pages

The user can choose to end the tutorial by clicking the “End Exam” button displayed in Figure 15. However, a warning message will be displayed to the user if they have not completed all of the questions (as displayed in Figure 16). This warning message is included to try and encourage the user to complete all of the tutorials and low stakes questions. If the user chooses to end the exam, they will be taken to a screen which summarises their performance within the tutorial (as shown in Figure 17). Included within this summary is the marks the user achieved on each of the tutorial questions as well as the option to “review” this specific tutorial and questions again to determine which answer they selected and whether it was correct or not.

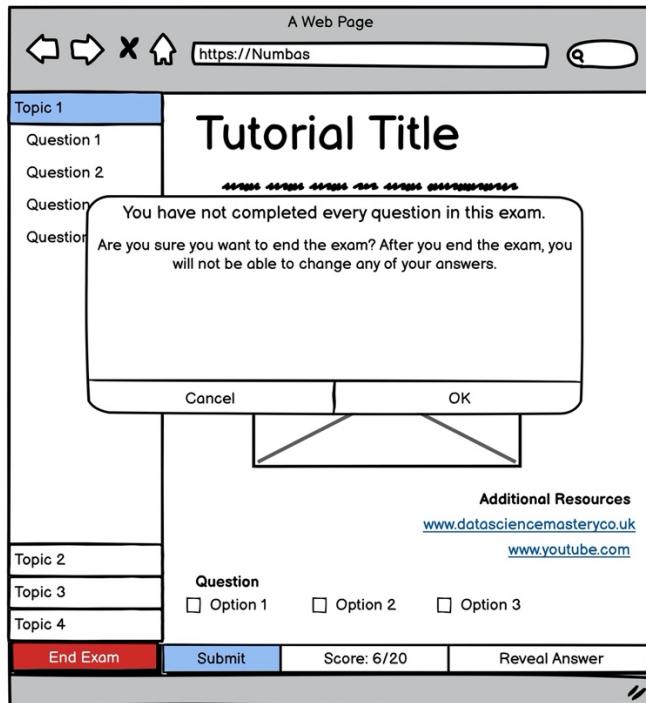


Figure 16: Message upon selection of 'End Exam'

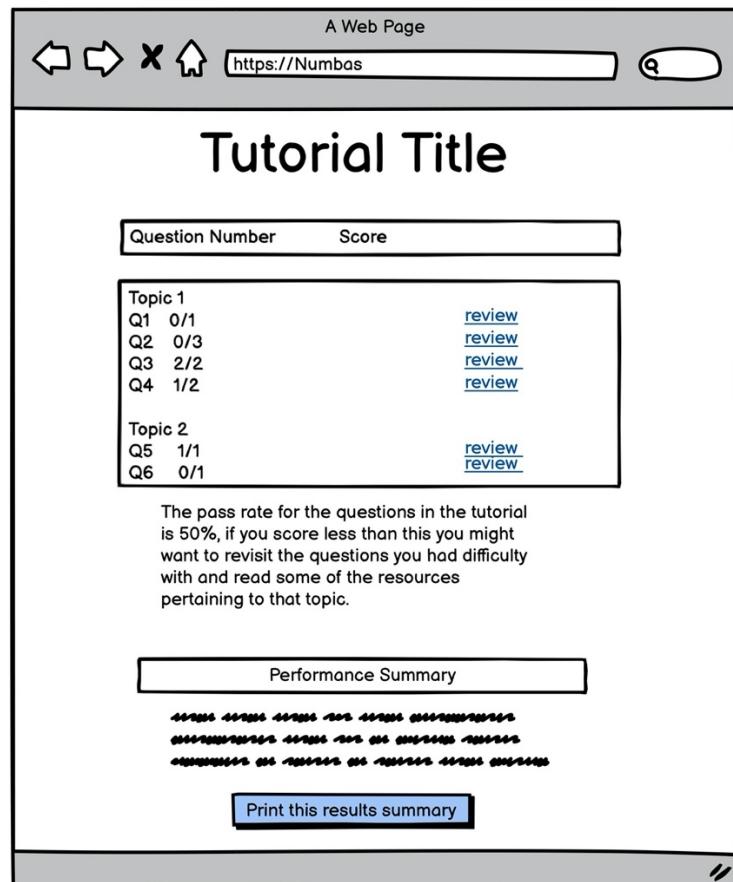


Figure 17: Performance summary

As the main framework of the online tool is already implemented through the existing Numbas system, the implementation phase is centrally focused on the content of the tutorials and the

importance of setting this at an appropriate educational level for a HE audience and ensuring it is engaging and informative.

5.6 Implementation

Creating the content for MetaLearning was a multi-faceted process which required not only the technical material to be created alongside the interventions aimed at confidence and self-efficacy. Throughout the instructional design process, the overall goal of the tool, to create an introductory course for Machine Learning, was kept at the forefront of all decision making. Determining the main demographics of the users of the tool was imperative to ensure the satisfaction of requirements R.4 and R.5. Adults enrolled within a university course were identified as the main users of the tool. The concept of adragogy and the core foundations of useful learning methods for adult learners were reviewed (Knowles, Holton. F and Swanson. A, 2014). Knowles *et al* (2014, p.4) advise that the six principles of adragogy are (1) the learner's need to know, (2) self-concept of the learners, (3) prior experience of the learner, (4) readiness to learn, (5) orientation to learning, and (6) motivation to learn. These six principles guided not only the content but the presentation of the material and relate to the requirements, specifically R.1 and R.5. Following the first and second principles of the learner's need to know and self-concept of the learners, the material within MetaLearning only covers content deemed essential for an introductory course for Machine Learning and also includes learning methods such as knowledge surveys and low stakes quizzes to help learners become more self-regulated, tying in with requirement R.2.

The students prior experience and readiness to learn were interpreted as the mental models students may have relating to their conceptions or misconceptions of Machine Learning and their previous educational background. Readiness to learn was also associated with the learning strategies and effective learning habits the users of the online tool may have. These principles were incorporated into the content through learning supports and clear explanations of the basics of the field. For example, AI, Machine Learning and Deep Learning all have introductory tutorials in which the content is based on the user having very little prior knowledge of these domains. Orientation to learning and motivation were encapsulated within the tutorials through real-life examples and scenarios, problem based learning and multi-modality content to keep user engagement high and to maintain user motivation following the requirements R.1, R.2 and R.5.

Based upon the findings from Chapter 4, the overarching content for the tool was decided. Covering AI, Machine Learning and Deep Learning was deemed necessary so that students gain comprehension of these areas and how they intersect. Within these sub-domains there will be tutorials introducing the field, providing a high-level explanation of some of the concepts and then further tutorials which go into detail of specific models and algorithms. Determining the best methods to frame this content and convey the often-complex ideas was guided by Oliver's (2001) framework for critical elements of online learning settings, shown in Table 17. The learning design elements of the framework include learning tasks, these are the activities and problems on which the learning is based and help to engage the learner in the material. For example, many of the tutorials within the online tool contain low stakes quizzes and real-life problems to help bed-in the knowledge gained from the content. The use of quizzes also corresponds to the learning supports category and requirements R.2 and R.3 as immediate feedback is provided to the user upon submitting their answer to the question.

Learning resources not only relates to the content, but additional material provided within the tutorials. As MetaLearning is an introductory resource to the field of Machine Learning, each tutorial will provide the users with a list of additional resources they can access which will provide further information at a higher knowledge level. Providing resources, from a variety of different perspectives and sources is important aspect of constructivist learning environments (Oliver, 2001).

Learning Design Elements	Description
Learning Tasks	The activities, problems, interactions used to engage the learners and on which learning is based.
Learning Resources	The content, information and resources with which the learners interact and upon which learning is based.
Learning Supports	The scaffolds, structures, encouragements, motivations, assistances and connections used to support learning.

Table 17: Framework for critical elements of online learning settings (altered from (Oliver, 2001))

One of the unique aspects of MetaLearning is the explicit focus on improving student's confidence within their learning of AI and the interventions in place to assist them in becoming more self-regulated learners. The interventions and strategies were determined based upon the findings from the initial study into students' experiences learning within this field and

educators' experiences teaching this subject (Chapter 4, p.81). The finding that both student and educators identify issues surrounding the theoretical aspect of AI led to a focus on mitigation strategies related to mathematics anxiety and comprehension level and the use of knowledge surveys to help users identify specific areas within mathematics and statistics which they are not confident with.

Aiming the content at absolute beginners, with no prior knowledge expected, alongside the use of low stakes questions specific to the content they have just learnt was aimed at bolstering student's self-efficacy and will fulfil requirements R.1 and R.2. Providing immediate feedback to the students was also important so that they can determine their progress and identify specific tutorials which they may need to revisit to improve their learning within this area, this will fulfil requirement R.3. Fulfilment of requirements R.4 and R.5 will require attention to detail and consideration of multi-modality content to ensure users have a consistent engaging experience using the tool. Upon completion of the implementation phase, thorough testing will need to be completed to ensure the satisfaction of requirements R.6 and R.7.

5.7 Tutorials

Implementation of the tutorials mainly consisted of tailoring the existing Numbas framework as outlined in Section 5.5 to ensure fulfilment of all of the requirements. Implementation of the tutorials and the content and inclusion of multi-modality learning materials were outlined in RO2.b (Chapter 1, p.6). Incorporation of strategies to improve student metacognition and self-regulation were also defined in RO2.c (Chapter 1, p.6).

In total there are six tutorials which comprise MetaLearning (links in Appendix I), each tutorial contains a range of 'questions' which contain individual lessons on a specific topic as outlined in requirement R.1. For example, the Overview of AI tutorial contains four 'questions' as displayed in Figure 18, each one contains material on a specific topic related to AI. The majority of 'questions' also contain a simple test at the end of the lesson to help students consolidate their learning as outlined in R.2 and R.3 and additional resources so they can undertake further work on a particular session if they wish.

NUMBAS

Overview of AI

Question 1	Score: 0/1 Unanswered	Artificial Intelligence (AI) is a field of computer science which endeavours to comprehend what intelligence is and to create intelligent machines.
Question 2	Score: 0/1 Unanswered	A few formal definitions of AI: Professor John McCarthy (coined the term AI):
Question 3	Score: 0/1 Unanswered	AI "is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable."
Question 4	Not marked	Oxford English Dictionary:
Total	0/3	AI is "the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages."

Pause

End Exam

There are many disciplines which contribute to AI including:



Figure 18: Overview of AI tutorial screen

Each of the six tutorials can either be uploaded to a VLE or be ran as standalone webpages, this ease of access is in fulfilment of requirement R.6. There is no specific order in which the tutorials are to be completed, if the user is a complete beginner it would be advisable to attempt the Overview of AI (Figure 19) and Machine Learning tutorials first to get a high-level understanding before delving into the other tutorials which are more in-depth and technical. However, the user can also choose to complete a specific tutorial if they think they need more learning within this area, for example with Maths and Stats for Machine Learning. If a user was wanting to specifically learn about Deep Learning, for example for a project, they could just complete the Overview of Deep Learning tutorial.

Overview of AI

Number of Questions: 4
Marks Available: 3
Pass Percentage: 50%

This tutorial will give you an overview of the field of artificial intelligence.
You may want to take notes throughout this tutorial.

This exam is running in standalone mode. Your answers and marks will not be saved!

Start

Figure 19: Introductory screen for the Overview of AI tutorial

Testing as a learning tool has been implemented within the tutorials as a cognitive intervention, this method offers both retesting and self-directed mastery opportunities which

have shown to reduce mathematics anxiety and improve confidence and self-efficacy as outlined in Chapter 2 (Sections 2.8.1 and 2.8.2, p.50). The use of testing as a learning tool as a potential strategy to improve metacognition and self-regulation was also outlined as RO2.c (p.6). Inclusion of this strategy will enable identification of the efficacy of this method as an intervention to overall improve aspects related to self-regulation, it will also enable insights to be drawn on student opinion of this technique as a learning tool. There is an overall pass mark for each tutorial which is set at 50%. If the student fails to reach the 50% target, they will receive feedback advising them to attempt the tutorial again. There is no limit to the number of times a user can access or complete these tutorials. When the user is first presented with the tutorial, they are advised on the pass percentage alongside the number of questions and marks available (as shown in Figure 19), this detail was added towards the fulfilment of requirement R.5 as quizzing has been linked with motivation (as outlined in Chapter 2, Section 2.9, p.55). The user is also advised to take notes throughout the tutorial, as findings from the case studies in Chapter 4 (Section 4.6.4, p.106) identified note taking as a common learning strategy used by students.

5.7.1 Maths and Stats for Machine Learning

A tutorial covering mathematics and statistics was created first. The decision to create a tutorial specific to mathematics and statistics for Machine Learning was informed by several findings from Chapter 4. The online review of modules (Section 4.3, p.82) identified that the majority of modules offered within the AI domain, specifically Machine Learning modules had a number of mathematics and statistics pre-requisites. In relation to this finding, variation in educational background was identified within the case studies (Sections 4.5.1 and 4.5.2, p.101). Also, lecturers participating in this research highlighted the difficulty that students have with mathematical aspects of their modules (Sections 4.4.1 and 4.4.2, p.97). Therefore, a tutorial offering instruction and guidance on this topic was formulated to eliminate some of these inequities in knowledge.

The inclusion of this subject in MetaLearning is unique as a lot of the existing courses do not cover this topic or only provide a brief recap which assumes the user already possesses a high level of mathematics knowledge as discussed in Chapter 4 (Section 4.6.1, p.123). Before starting the lessons, the user is prompted to complete a knowledge survey. A knowledge survey is a method to assess progress in learning and intellectual advancement, it requires

students to respond to specific content questions with an evaluation of their confidence in being able to competently answer the content question (Nuhfer and Knipp, 2003). The knowledge survey is repeated at the end of the tutorial so that students can reflect on how their confidence has increased as a consequence of their learning, this awareness can lead to an increase in self-efficacy which can encourage further self-regulation (Nilson, 2013b). Inclusion of this strategy pertains to RO2.c (Chapter 1, p.6).

A major influence on the choice of topics to be covered within this tutorial was the varying educational backgrounds pertaining to mathematics qualifications as outlined in Chapter 4 (Sections 4.5.1 and 4.5.2, p.103). The results from the case studies indicated a wide spread of mathematics attainment from students, with some indicating that their highest qualification is GCSE Mathematics, to some students whose undergraduate degree was in Mathematics and Statistics. Due to the variation, the decision to aim the material at a beginner's level was determined and informed the content of this tutorial. As the tutorials are optional, students with a high level of mathematics skill would not be required to complete this tutorial, this decision was informed by requirement R.5 as users may lose motivation to continue using the tool if they feel the content is not beneficial to their learning. However, it may provide a comprehensive revision session for topics specific to understanding the theoretical basis of many Machine Learning aspects for students who possess a high mathematics skill level.

The material for the tutorial was split into five main categories, Maths, Linear Algebra, Differential Calculus, Probability and Statistics. These categories were based on the module pre-requisites identified through the online review of modules in Chapter 4 (Section 4.3, p.82). Table 18 details the question groups and an overview of the content within each group.

The topics selected for inclusion within this tutorial were also selected based on their foundational importance in understanding the inner workings of many of the algorithms within the Machine Learning domain. Deisenroth, Faisal and Ong (2020) outline a number of foundational topics including linear algebra, matrix operations, probability theory and calculus, all topics which are covered in the Maths and Stats tutorial. An example of the relevance of these topics is how inherent for example, linear algebra is in many areas of Machine Learning including algorithms like Principal Component Analysis (PCA) (Jolliffe and

Cadima, 2016) as well its importance in understanding how optimization works (Akinfaderin, 2017).

Question Group	Content
Maths	Scalars, vectors and matrices, matrix operations, exponents and logarithms, summation and product.
Linear Algebra	Multivariable functions, linear and non-linear functions, linear equations and transformations.
Differential Calculus	Derivatives, derivative rules, partial derivatives.
Probability	Probability theory, Bayes theorem, probability distributions.
Statistics	Correlation, curve fitting, regression, standard deviation.

Table 18: The question groups and overview of content of the Maths and Stats tutorial

The level of mathematics required depends on the project or course, it is not always essential to have an expert level of mathematics comprehension. For example, there are a number of libraries and packages such as Weka (University of Waikato, 2020) and Keras (Chollet, 2015), which do not require the user to have in-depth theoretical knowledge to be able to build and apply Machine Learning models. However, having an understanding of some of the mathematical theory underpinning Machine Learning concepts can help the student to make better informed choices when choosing the model, choosing hyperparameters and validation settings as well as identifying underfitting and overfitting (Akinfaderin, 2017). Choosing an appropriate model for specific datasets was discussed as an issue which students encounter in Chapter 4 (Section 4.6.2, p.124).

The inclusion of low-stakes quizzing enables the user to consolidate their knowledge in the content they have just read (Sotola and Crede, 2021). There are a number of question types included in this tutorial which are supported by Numbas, such as multiple choice, several step number entry and matrix entry. Multiple choice questions have been shown to “trigger beneficial retrieval processes for both tested and related information” especially if the choices include competitive incorrect alternatives (Little and Bjork, 2015). The number entry questions require the user to answer the question in a number of steps to encourage them to consider the problem-solving process and to decompose the problem into smaller subtasks. The matrix entry questions pertain specifically to the content relating to matrices which require this format for accurate representation of the solution. The variation in question types

was deliberate to try and maintain engagement in the content as successive multiple-choice questions may be laborious for some users, this aligns with requirements R.1 and R.5. Feedback is provided for both individual questions and for the overall tutorial to fulfil requirement R.2, as shown in Figure 17. The summary of marks displayed upon tutorial completion (shown in Figure 20) also provides a breakdown of specific questions the user got wrong to aid them in identifying specific topics which require further study. This summary page also displays the percentage of questions answered correctly and guides the user to revisit questions they had difficulty with and to utilise the additional resources included in each question page if they scored less than 50%.

Probability			
Question 17	1	/	1
Review			
Question 18	1	/	1
Review			
Question 19	1	/	1
Review			
Question 20	1	/	1
Review			
Question 21	0	/	1
Review			
Statistics			
Question 22	1	/	1
Review			
Question 23	0	/	0
Review			
Question 24	1	/	1
Review			
Question 25	0	/	0
Review			
Question 26	1	/	1
Review			
Post Maths and Stats Knowledge Survey			
Question 27	0	/	0
Review			
Total	15	/	26 (57%)
<p>The pass rate for the questions in the tutorial is 50%, if you score less than this you might want to revisit the questions you had difficulty with and read some of the resources pertaining to that topic.</p> <p>Thank you for using this tool, in order to improve the system please complete the questionnaire linked below:</p> <p>https://forms.ncl.ac.uk/view.php?id=6719176</p>			

Figure 20: Summary of marks from the Maths and Stats tutorial

5.7.2 Data Pre-Processing

One of the most fundamental stages of a Machine Learning project is data processing as it has been shown to have critical influence on the model performance (Huang, Li and Xie, 2015). The starting point for this topic was an explanation of data and the different types of variables. The Data question defines the term ‘variable’ and explains the difference between discrete,

continuous and categorical variables. This question concludes with an example variable and the user must decide whether this is continuous or discrete (as shown in Figure 21).

Would the following be a discrete or continuous variable:

The length of time it takes a truck driver to drive from Sunderland to London

Discrete Continuous

[Submit answer](#) Score: 0/1 [Try another question like this one](#) [Reveal answers](#)

Figure 21: Example of a multiple-choice question relating to variables

A question included within this tutorial, The Preparing the Data question covers three main stages related to processing the data for input into a Machine Learning model, these include selecting the data, pre-processing the data and feature engineering. Feature engineering, a term widely used within Machine Learning, relates to “the process of formulating the most appropriate features given the data, the model, and the task” (Zheng and Casari, 2018, p.3). Within the case studies, the students who completed the one-minute paper identified issues comprehending feature engineering (Chapter 4, Section 4.5.1, p.113) therefore inclusion of this topic within MetaLearning was important to provide a resource for learning. Specific aspects of feature engineering covered within this tutorial include representation transformation, feature tuning, extraction, decomposition and aggregation.

Data representation, the process of mapping data to useful features is also covered within the Data Pre-Processing tutorial. Different mappings are demonstrated such as numerical representation, categorical to numerical and numerical to categorical. The user is presented with an example of a mapping of a dataset into a set of features and asked to determine what type of mapping has been undertaken as shown in Figure 22.

Question 2	Score: 0/1 Unanswered	If we had: age{1,5, 7, 18, 54, 25} and transformed it into {'infant', 'child', 'child', 'adult', 'adult','adult'}
Question 3	Not marked	What type of mapping have we undertaken?
Question 4	Score: 0/1 Unanswered	<input type="radio"/> Categorical to numerical <input type="radio"/> Numerical representation <input type="radio"/> Numerical to categorical
Total	0/3	

[Display options](#) [Submit answer](#) Score: 0/1 [Try another question like this one](#) [Reveal answers](#)

Figure 22: Low stakes question pertaining to mapping of a dataset

The purpose of this question was to demonstrate and visualise what the mapping process looks like alongside asking the user to reflect on what they had just learnt to determine what type of mapping had occurred.

5.7.3 Overview of AI

Including content related to AI was essential for the user to gain cognition of how Machine Learning and Deep Learning are situated within this field. The first page within the Overview of AI tutorial provides a definition of AI and the many disciplines which contribute to the field such as Computing, Philosophy, Mathematics and Neuroscience (shown in Figure 23). The aim of this visualisation was to demonstrate to the learners the interdisciplinarity of the field and to enable them to consider the type of skills they will need to become competent practitioners within this area, potentially assisting them in building a more accurate mental model.

NUMBAS	
Overview of AI	
Question 1	Score: 0/1 Unanswered
Question 2	Score: 0/1 Unanswered
Question 3	Score: 0/1 Unanswered
Question 4	Not marked
Total	0/3
Display options	

Encyclopedia Britannica:
"AI, the ability of a digital computer or a computer-controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience."

There are many disciplines which contribute to AI including:

```

graph TD
    Computing((Computing)) --> AI[AI]
    Philosophy((Philosophy)) --> AI
    Mathematics((Mathematics)) --> AI
    Neuroscience((Neuroscience)) --> AI
    Psychology((Psychology)) --> AI
  
```

Figure 23: Visualisation of the disciplines which contribute to the field of AI

Visualising how the three fields intersect and the relationship between Machine Learning and Deep Learning provides a basis of knowledge upon which the tutorials specific to these areas can build upon (as shown in Figure 24). This visualisation and overarching content was included to align with instructional design principles, specifically scaffolding users to level 5 of Bloom's taxonomy, *synthesis*, (as outlined in Chapter 2, Section 2.5.4, p.36) where learners extend their understanding beyond what they have been instructed on. This is the aim of this tutorial, to provide learners with a foundational understanding and mental model of the domain, in which they can then build upon with more in-depth knowledge on the specific subdomains of AI, specifically Machine Learning.

NUMBAS

Overview of AI

Question 1	Score: 0/1 Unanswered
Question 2	Score: 0/1 Unanswered
Question 3	Score: 0/1 Unanswered
Question 4	Not marked
Total	0/3

Within the field of AI there are many subdisciplines such as machine learning, deep learning and reinforcement learning which will be covered in later tutorials.

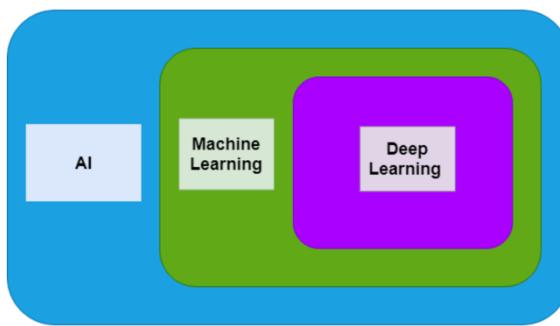


Figure 24: Visualisation of the differing subdomains within AI and how these intersect

A history of the field was provided to draw attention to some of the milestones and major breakthroughs in AI, this was used to illustrate the progress that has been made within this domain and how it has led to the current state of the art. This content was presented as a visualisation of a timeline to engage the viewer and to make it easier to present a large amount of content in an appealing way (Appendix M). There is also a tutorial on the Turing Test, originally titled the imitation game (Turing, 1950) as this is highly influential within the field of AI and many of the users may be familiar with Alan Turing but not fully understand his impact and contribution.

Inclusion of content associated with the ethical considerations of AI was incredibly important to make students of this domain aware of some of the issues within this field and to set them on the path to becoming responsible practitioners. Alongside real-life examples of unethical AI applications and instances when the technology has proven less than beneficial, a selection of considerations are also discussed such as bias, transparency, security and equality (see Section 2.3, p.19).

5.7.4 Overview of Machine Learning

Before delving into specific algorithms, it was important to cover the basics, such as defining Machine Learning, explaining what learning algorithms are and the different types of learning. The types of learning algorithms covered include supervised, unsupervised and reinforcement learning. Also included in this definition was a graphic which acts as a quick reference guide to create distinctions between the groups (as shown in Figure 25). Including specific examples of algorithms which fall into the differing types of Machine Learning was a way to initially familiarise the students with models they would be looking at in further tutorials.

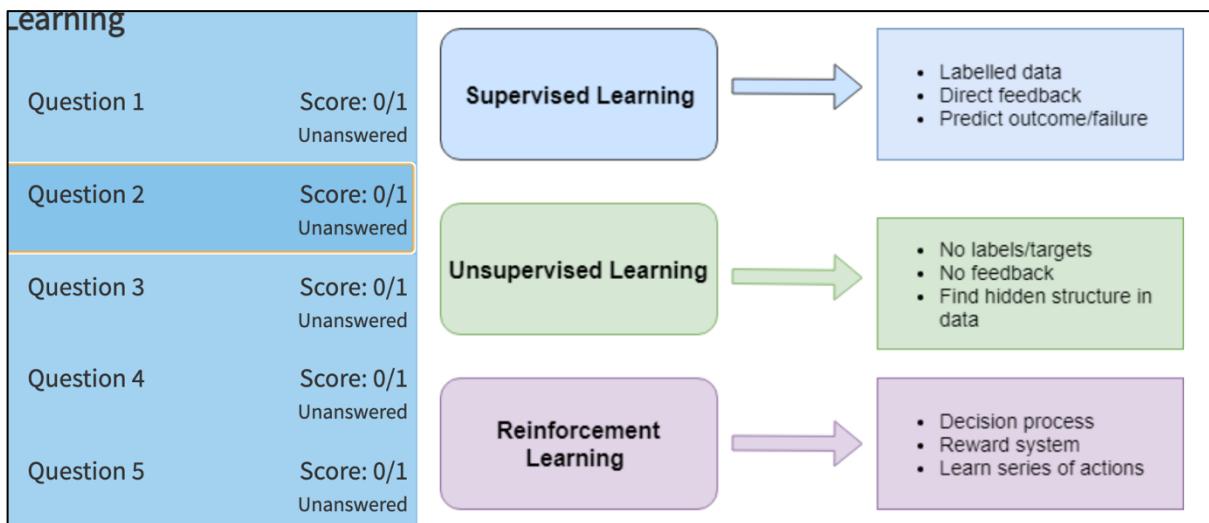


Figure 25: A visual reference guide to the different types of learning algorithms

Covering the issues of underfitting and overfitting introduced the users to keys issues within Machine Learning such as generalization and the bias/variance trade off. Another key issue addressed within this tutorial is the Machine Learning workflow, this clearly delineates the different stages of a Machine Learning project. Within the workflow key concepts were introduced such as optimization and activation and loss functions. Acquainting users with these concepts within the workflow enabled these terms to be explained within the context of a Machine Learning project and hopefully removed ambiguity surrounding when and for what purpose these are implemented.

All of the tutorials in the Overview of Machine Learning have questions that help with comprehension of the material as it is pivotal to the other tutorials, particularly Machine Learning Algorithms and Overview of Deep Learning. Question types include true/false and multiple choice based on real-life examples. Inclusion of real-life examples is important to engage the users in how this technology is currently being employed as well as spark consideration of how the users can start applying Machine Learning for their own projects.

5.7.5 Machine Learning Algorithms

The Machine Learning Algorithms tutorial aims to give users an understanding of specific algorithms within supervised and unsupervised learning, reinforcement learning and evolutionary algorithms. Deciding which algorithms to cover within these categories was informed by common algorithms within these sub-domains as well as the findings from the

online review of modules and lecturer questionnaires in Chapter 4 (Sections 4.3.2 and 4.4.1) which identified common topics taught on traditional AI courses.

This tutorial is split into four sections for each type of learning, starting with supervised learning. Algorithms covered in this section include logistic and linear regression (Worster, Fan and Ismaila, 2007), Naïve Bayes classifier (Gandhi, 2018), Decision Trees (Quinlan, 1986), Random Forest (Ho, 1995) and Support Vector Machines (Cortes and Vapnik, 1995). To try and assist the students with their learning, where possible, a visual explanation was included which aimed to explain the model within a different modality. For example, within the Decision Trees question, two graphics were created, one to explain the hierarchical design (displayed in Figure 26) and the other to explain information gain.

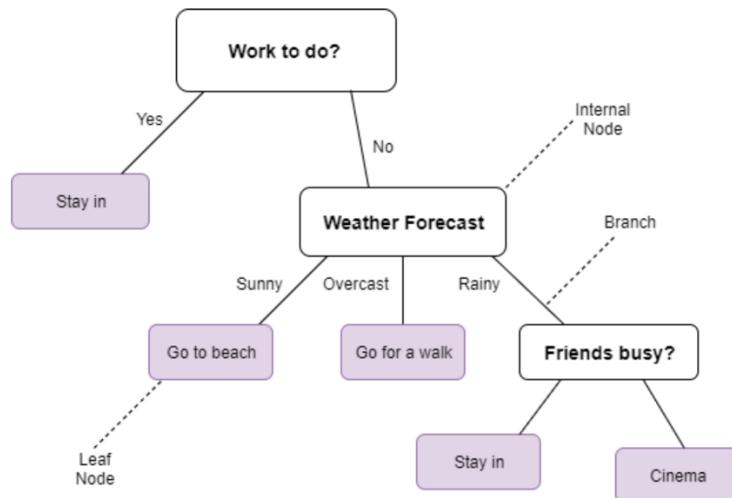


Figure 26: Visual representation of a Decision Tree

The majority of the tutorials also contain code snippets or examples and include additional resources to ensure that the user has access to further information if required. Multiple choice questions were used to enable the user to reflect on the content they have just read and to give immediate feedback. The user can then use this feedback to either revise the specific model again or move onto the next section of the tutorial.

Within the Unsupervised Learning section two different ways of applying these algorithms were covered, the use of clustering to discover relationships within a dataset and dimensionality reduction, linking back to feature extraction covered within the Data Pre-processing tutorial. The K-Means algorithm (Bock, 2008) was the clustering technique chosen.

To accompany the written descriptor, a visualisation of the step-by-step process involved in implementing this algorithm was created as shown in Figure 27.

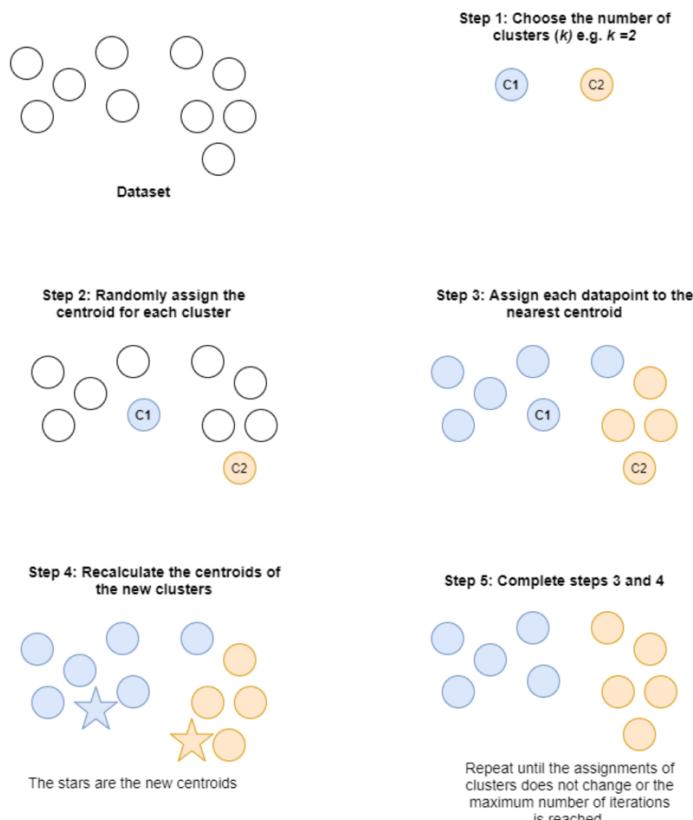


Figure 27: Visualisation of the steps for implementing the K-Means algorithm

Within the Reinforcement Learning section, further instruction was given on what reinforcement learning is and how it works, the terminology used within this learning domain and the different approaches. The Q-Learning algorithm (Watkins and Dayan, 1992) was detailed as a specific type of this kind of learning and the steps for implementing this approach were explained. Evolutionary algorithms were also included in the Machine Learning Algorithms tutorial and within this category genetic algorithms (Katoch, Chauhan and Kumar, 2020) were discussed. To assist learners with their comprehension of genetic algorithms, key terms were explained with visual representation and inclusion of pseudo code. To summarise genetic algorithms, the advantages and disadvantages were detailed, and additional resources were provided if further tuition was required.

5.7.6 Overview of Deep Learning

The advances within Deep Learning have led to an increase in uptake of this technology within industry and as a consequence, a growing interest amongst students to become proficient within this domain. The content for this tutorial was determined based upon the findings

outlined in Chapter 4 from the online review of modules (Section 4.3, p.82) and from the interviews with the lecturers (Section 4.4.2, p.97). These findings highlighted key areas of Deep Learning such as Convolutional Neural Networks (LeCun *et al.*, 1998) and Recurrent Neural Networks (Hochreiter and Ugen Schmidhuber, 1997) as well as aspects of Deep Learning identified as troublesome such as backpropagation (Rumelhart, Hinton and Williams, 1986).

As Deep Learning was identified as a potential threshold concept, situating one of the main aspects of this domain, Artificial Neural Networks (Abiodun *et al.*, 2018) within the context of human learning was chosen as a method to help users build a stronger mental model. Describing how Artificial Neural Networks work alongside the concept of a neuron (both artificial and biological) in conjunction with how humans learn enabled a range of concepts to be explained such as backpropagation (Rumelhart, Hinton and Williams, 1986). However, it was important to reiterate that AI research has a lot of progress to make before we can realistically model human learning. The aim of the inclusion of human learning was to ease the users into the subject of Deep Learning within a context the student understands as well as promote active engagement as students may be enthusiastic about learning about themselves.

The inclusion of the popular Deep Learning models, Convolutional Neural Networks and Recurrent Neural Networks within this tutorial, required addition of material on backpropagation and gradient descent as these are key for training. Within the gradient descent tutorial, a graphical representation and analogy of a person climbing down a hill was used to aid the users understanding of the process and the purpose of this algorithm, as shown in Figure 28.

An easy way to think of gradient descent is as a person climbing down a hill, we want to climb down until we reach some minimum cost.

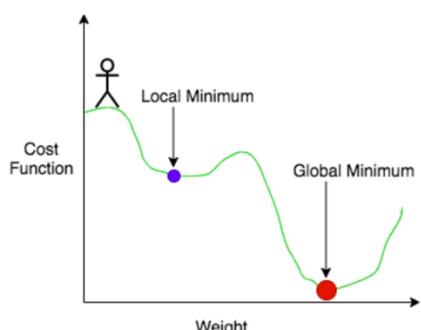


Figure 28: Graphical representation of gradient descent

Deciding how much detail to include in the backpropagation tutorial was an issue to consider as this algorithm can be difficult to understand. However, as the tutorials are aimed at beginners to this field it was decided to keep the explanation of backpropagation to a high level so that students would have an understanding of its purpose and loosely how it works, without explaining the in-depth details.

The purpose of the inclusion of Generative Adversarial Networks (GANs) (Goodfellow *et al.*, 2020) was to enable the user to gain experience of a wide range of models. GANs were also covered on the Deep Learning modules identified in the online review of modules (Chapter 4, Section 4.3.3, p.91). Overall, the Deep Learning tutorial covers five different models, where possible the material includes visualisations to aid the user in their understanding. The use of multiple-choice questions was included to try and boost the users confidence in their understanding of the material they have just studied.

5.8 Testing

Upon completion of the implementation phase, MetaLearning was reviewed by a Senior Lecturer in Data Science and Machine Learning to ensure the content was accurate and pitched at an appropriate level for a HE audience. Any feedback pertaining to the material was implemented to ensure the quality of learning material was at the highest possible standard.

As Numbas is an extant system, this software has been through extensive usability testing as a good user experience is a central tenet of the aim of Numbas (Perfect, 2015). Numbas has also been used at a variety of differing institutions such as University of Leicester and Royal Darwin Hospital for educational purposes (Newcastle University, 2018). Therefore, further usability testing was not deemed necessary at this stage.

The requirements set out in Section 5.3 were at the forefront of the design and implementation phase of the tutorials and were reviewed within the testing stage to ensure they had been fulfilled within the system. The following sections outline the extent to which these requirements have been satisfied within the online tool.

5.8.1 R.1 Ability to host engaging multi-modality content

Requirement R.1 centred around the content, ensuring it was pitched at an appropriate educational level and stimulated and held the users attention in the subject matter. This requirement was fulfilled through inclusion of a number of different aspects and features implemented during the design and implementation phase. All content included within the tutorial was informed by the findings from Chapter 4 (p.81). The online review of modules, questionnaires and interviews with the lecturers and case studies enabled identification of domain content taught within the HE sector as well as topics which were deemed pivotal and potential threshold concepts. Using these findings as a benchmark for the online tool content ensured that the material was pitched at the right educational level.

Requirement R.1 was also satisfied in relation to the multi-modality aspect of this requirement. As evidenced by examples such as Figures 27 and 28, a number of concepts were given a graphical representation, especially for complex ideas such as gradient descent. Representing a complex idea as something the students can already relate to and already have a mental model of was a strategy to enable them to understand complicated ideas within this domain. Alongside visualisation, the idea of using sonification to assist in the explanation of difficult topics was also researched, however inclusion of this strategy was not deemed essential at this stage and will be investigated at a later stage.

The use of case studies was also a method employed to engage users with the content, although this technique could have been extended further it was used within the Recurrent Neural Networks tutorial for example to show how a chatbot utilises this algorithm. Real life examples were also employed, for example the idea of ANNs was explained in conjunction with how humans learn, so that the users could correlate their understanding of how they learn with that of learning algorithms.

As well as engaging the users with the content, it was also important that the tutorials ultimately supported student learning and the progression from surface to educational deep learning. The tutorials support the concept of surface learning (as outlined in Chapter 2, Section 2.5.3, p.32) as the learners are familiarised with the subject matter vocabulary and foundational content. Strategies for surface learning are encouraged by the tool including note taking and imagery. Practice testing, through the use of low stakes quizzing as

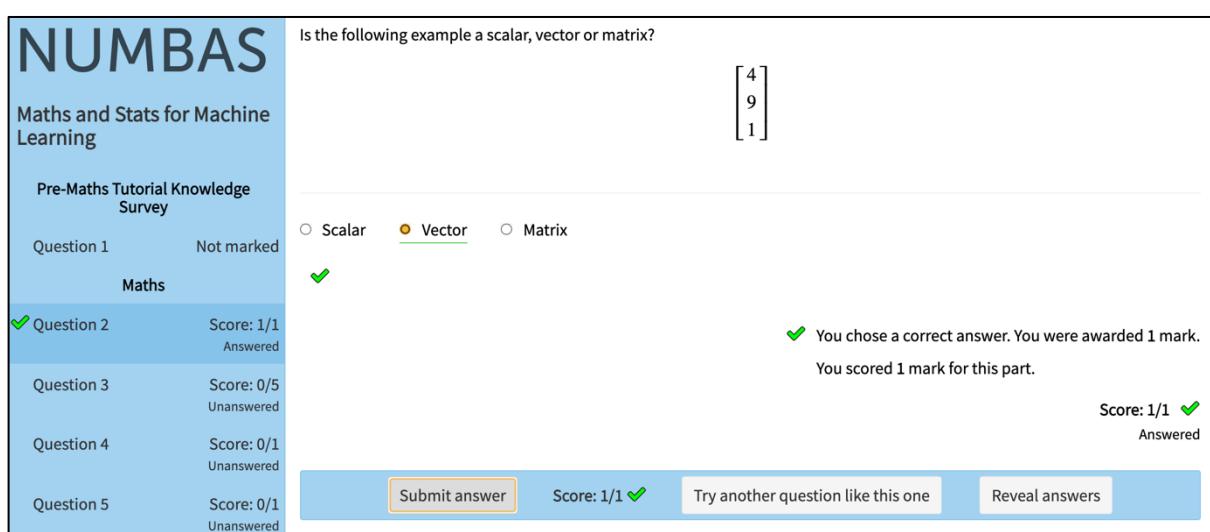
recommended by Hattie and Donoghue (2016) helps students through the consolidation phase of surface learning, where learners encode this new information so that it can be retrieved at a later date. After the consolidation phase, students potentially develop educational deep learning, supporting students into this learning phase is facilitated in MetaLearning through the knowledge surveys and low stakes quizzing as this enables self-monitoring and questioning.

5.8.2 R.2 Capability for formative assessment

This requirement is fulfilled as the inclusion of testing as a learning approach has been implemented within each tutorial through low stakes questioning. To engage the users with this learning strategy, alternative modes of questioning were created including multiple choice, matrix entry boxes and text boxes. Upon submission of the answer to the question, the user receives immediate feedback on whether they answered correctly or not. Upon completion of the exam the users also get a summary of how many questions they got right overall. An important aspect of testing as a learning approach is the option of retesting, this was implemented within MetaLearning and the users can retake the questions and revisit the tutorials as many times as they wish.

5.8.3 R.3 Provide learners with feedback

The fulfilment of this requirement was essential to ensure the testing as a learning tool approach enabled learners to identify specific topics in which they required further learning. The most basic form of feedback which was implemented was acknowledgement of whether the user had correctly answered the low stakes question as shown in Figure 29.



The screenshot shows a NUMBAS interface for a 'Pre-Maths Tutorial Knowledge Survey'. The question is: 'Is the following example a scalar, vector or matrix?' The example is a column vector:
$$\begin{bmatrix} 4 \\ 9 \\ 1 \end{bmatrix}$$
. The user has selected 'Vector' as the answer, which is correct. The interface shows the following details:

- Question 1:** Not marked, Maths, Scalar (radio button)
- Question 2:** Score: 1/1, Answered, Vector (radio button, selected), ✓ (green checkmark)
- Question 3:** Score: 0/5, Unanswered
- Question 4:** Score: 0/1, Unanswered
- Question 5:** Score: 0/1, Unanswered

At the bottom, there are buttons for 'Submit answer', 'Score: 1/1 ✓', 'Try another question like this one', and 'Reveal answers'.

Figure 29: Example of feedback of a correct answer

One of the advantages of this form of learning strategy is that the user can retake and submit the answer as many times as they wish to enable them to correctly answer the question, leading to a successful mastery experience.

As well as providing feedback on whether the answer was correct, the user can also select the 'Reveal answers' button and receive additional feedback on the expected answer, as displayed in Figure 30. This form of feedback allows the user to identify and reflect on the differences between their incorrect answer and the expected answer.

Each section of the multiplication is broken down to help you.

a)

$$4 \times 5 + 6 \times 4 + 3 \times 2 =$$

50 ✓

Expected answer: 50

✓ Your answer is correct. You were awarded 1 mark.
You scored 1 mark for this part.

Score: 1/1 ✓
Answered

b)

$$4 \times 3 + 6 \times 8 + 3 \times 5 =$$

73 ✗

Expected answer: 75

Figure 30: Example of feedback with an advice section

Additional resources are also provided within each tutorial to ensure that the learner has opportunity for further learning if they are struggling with a particular topic. Additional resources for some of the tutorials, particularly the ones relating to Machine Learning are links to code examples where the learner can utilise these examples to embed their theoretical understanding within a practical context.

As well as feedback on individual questions, the user is also given overall feedback on their mark for the full tutorial, Figure 20 demonstrates this. This form of feedback allows the learner to reflect on their understanding of the topic and the tool prompts them to revisit the content and additional resources if they scored below 50%.

5.8.4 R.4 User interface must be attractive, consistent and responsive to the user

As the online tool used the existing Numbas framework, changes to the user interface were limited. However, as Numbas has already been subject to thorough usability testing, requirement R.4 was fulfilled. The visual design and navigation of the tool was consistent throughout each tutorial and individual question pages. Regarding accessibility, the user had the option to customise and change display options as shown in Figure 31. However, the accessibility features within Numbas should be extended to allow for the input of alternative text for images and to ensure full compliance with the use of screen readers.

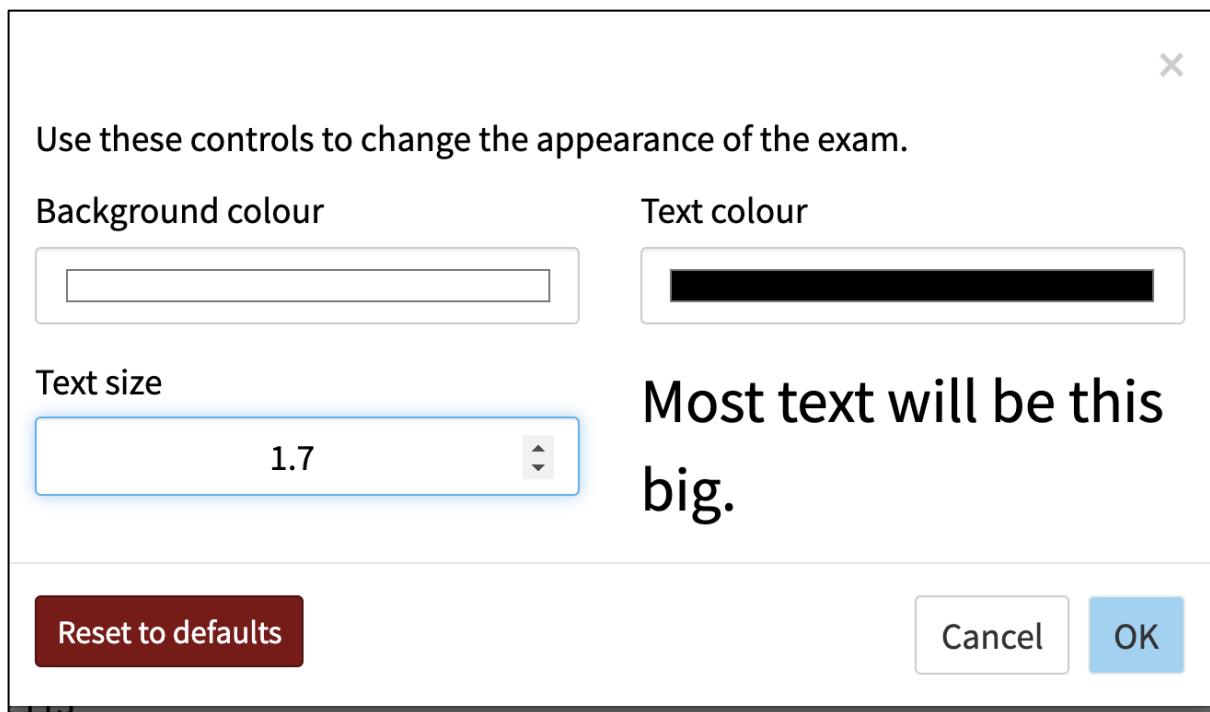


Figure 31: Options to change the visual display

5.8.5 R.5 Tutorials must be interactive and engage users to maintain motivation

Instructional design principles outlined in Section 5.3.2 guided the implementation phase of development to ensure fulfilment of this requirement. Although full satisfaction of this requirement cannot be assessed until users provide feedback on their experience using the tool, varying principles have been included. The tutorials have an element of interactivity in that they contain low stakes questions for the students to complete as well as providing personalised feedback on their score in the quizzes. To maintain motivation, the users can resubmit their answers to the questions as many times as they wish so that they can see a successful outcome to the tutorial. Additional resources are also provided to enable users to further their learning beyond the scope of MetaLearning.

5.8.6 R.6 Ensure access to the tool is readily available

Requirement R.6 was fulfilled as MetaLearning is accessible through two main methods. The tool can be integrated into a VLE, compatible VLE's include Blackboard (2021), Canvas (Instructure, 2021) and Moodle (2021). The main advantage of integration into a VLE is that the users scores and attempt data can be stored and retrieved, this is especially useful if the tool is being deployed to investigate the student's current level of understanding. However, users must be notified that this data is being collected. MetaLearning is also accessible as a standalone web page, Numbas and therefore MetaLearning is compatible with all major browsers and devices. When the tool is being ran as a standalone web page, the scores and data pertaining to the number of attempts or time spent on the tutorial is not stored.

5.8.7 R.7 No new personally identifiable data should be collected through the tool

Requirement R.7 has been realised as the methods for accessing the tool, outlined in requirement R.6 ensure that the user does not need to sign up or create a profile to access MetaLearning. When the tool is accessed through integration into a VLE, data pertaining to the users score and learning analytics such as the time spent on each tutorial is recorded. However, apart from the user the only people who have access to this data are individuals who already have access to that specific users VLE, for example a module leader.

When using MetaLearning as a standalone web page, none of the data created through use of the tool is stored or accessible when the web page is closed. Satisfaction of this aspect of the requirement was an important consideration pertaining to improvement of learners confidence so that they know that any answers to the low stakes questions will not be shared or available to others.

Ensuring that the answers to the knowledge survey contained in the Maths and Stats tutorial were anonymous irrespective of how the tool was accessed and signposting this to the user was incredibly important so that users could be honest when answering the questions and therefore being able to fully reap the benefits of this mitigation strategy.

5.8.8 Expert Requirement Verification

To independently verify the extent to which the requirements for the MetaLearning tutorials were fulfilled, a selection of the tutorials were reviewed by an expert in AI who currently teaches and is an active researcher within this topic at a post 1992 university. The verification process consisted of the expert evaluating whether they felt the requirements had been met and any comments they wished to document.

The expert determined that all requirements had been met within the MetaLearning tutorials. However, there was scope for some improvement pertaining to R.1 and R.5 advising that they liked the mix of images, text and question types, however, use of short videos could be included to potentially reduce cognitive load. Details of further iterations to MetaLearning will be discussed in Section 8.5.

5.9 Example of Use

Alongside the fulfilment of the requirements, it was important to ensure the flexibility of use for MetaLearning to meet the needs of differing educational scenarios as outlined in Section 5.4. One of the ways in which the tool can be employed is within blended learning, for example within a lesson covering Machine Learning. Both the lecturer and students will interact with the tool for differing purposes, the lecturer will use the tool as a resource for instruction and further engagement with the concepts covered within the session. The students will use the tool to build upon their knowledge within a specific area, the use of the tutorials will also give the students the opportunity to enquire and discuss particular models or questions with the lecturer and engage in discussion with other students.

For use within a lesson covering the basics of Machine Learning, MetaLearning can be included within the lesson plan as part of an activity. For example, the lesson may start with an introductory lecture on the basics of Machine Learning where concepts may be covered such as supervised and unsupervised learning, application examples and a basic guide to a model such as logistic regression. The students would then be required to complete the Machine Learning Algorithms tutorial as part of the activity section of the lesson plan. Whilst completing the tutorial, the students will encounter topics which they have previously discussed in the lecture session of the lesson, this will help solidify the new concepts as well as contextualise them within specific Machine Learning models. The students will also

encounter new concepts and will gain an understanding of how well they have understood these through the low stakes quizzes. As the students are using the online tool within the supported environment of a lesson, they can discuss any questions they have with their peers and with the lecturer.

Upon completion of the activity section of the lesson, the lecturer could either facilitate a discussion on the tutorial, for example by giving the students some real-life examples and asking which algorithm they would use for this problem and why. The lecturer could also choose a specific algorithm covered within the tutorial and go through a step-by-step guide of how to implement it for example, within Python.

The online tutorials could also be used as pre-work before the next lesson, for example if the lecturer had prepared to cover Deep Learning within their next session, they may ask the students to complete the Overview of Deep Learning tutorial to prepare and develop a basis of understanding which the lecturer will build upon within their lecture.

5.10 Summary

This chapter describes the online tool for Machine Learning called MetaLearning, a set of tutorials created as an introductory course to topics related to AI, Machine Learning and Deep Learning. Alongside the content a variety of mitigation strategies have been implemented within the tool such as knowledge surveys to ease mathematics anxiety, low stakes quizzing and timely feedback to increase student self-efficacy, which in turn can impact student self-regulation. Referring back to Table 17, the framework for critical elements of online learning (Oliver, 2001), all three categories of learning design fundamentals have been interwoven throughout MetaLearning. Learning tasks are prevalent throughout the tutorials due to the inclusion of low stakes quizzing as a learning tool as well as prompts for the learners to use learning strategies such as note taking and retesting. Relating to learning resources, the findings from Chapter 4 informed all of the content to ensure it was appropriate for students at a HE level and could be used as a complete introduction to the AI domain, with particular focus on Machine Learning. Learning supports and scaffolding were included through the use of visualisation of complex topics, the option for retesting, and the creation of content without any pre-requisite of prior knowledge within this domain.

The content and layout of the tutorials were designed to be engaging for the user, through multi-modality material and real-life examples. The ease of deployment of the tool was imperative to ensure ease of access for both educators and students. The option of anonymity as a standalone resource was important to ensure the user felt comfortable answering the questions and that they felt they weren't being assessed or judged on their progress.

The next chapter reviews the deployment of the tool and feedback from users, both students and lecturers, based on their experience using the technology, what content they found the most difficult and whether they felt the learning strategies were effective.

Chapter 6. Evaluation of MetaLearning

6.1 Introduction

MetaLearning was designed to serve as an introductory course for AI, specifically Machine Learning and Deep Learning. Alongside the tutorials, mitigation strategies were included to ease mathematics anxiety and to improve student self-efficacy. To determine how effective the online tool was at the stated objectives, it was trialled across two institutions with both lecturers and students.

Two of the institutions which evaluated MetaLearning were from the case studies outlined in Chapter 4 (p.101), University A and University B, although with different student cohorts. All participants who used the tool were not obligated to provide feedback; this was optional. Students were asked to provide feedback through an online questionnaire, lecturers provided their opinion on the tool through an interview.

In September 2021, a workshop was held at the United Kingdom and Ireland Computing Education Research (UKICER) conference entitled “Measuring the Difference Between Student and Staff Perception of Self-Efficacy and Confidence Using Online Tools” (United Kingdom and Ireland Computing Education Research, 2021) alongside colleague Laura Heels also from Newcastle University (ResearchGate, 2021). Within this workshop, participants were asked to review two of the MetaLearning tutorials, particularly in respect to the integration of strategies to improve learner confidence and self-efficacy. The following sections outline the specific data collection methods used, the analysis methods and the findings.

6.2 Comparison with Existing Tools

Determining how MetaLearning compares with other existing tools providing an introduction to AI enables insight into the unique aspects of MetaLearning. Comparison has been made with a range of tools which have previously been discussed within this research including the AI and Machine Learning module hosted by Code.org (2021a), as discussed in Section 2.6.2 (p.41). Specific courses created by the MOOCs outlined within Section 5.2, including courses offered by Udacity (2020) and Coursera (2021) have also been reviewed against the offering

provided by MetaLearning. Table 19 displays a comparison of key features of these educational online tools.

Online Course	Introductory Mathematics Content?	Study Skills Support?	Feedback Provision?	Multi-Modality Content?
MetaLearning	Y	Y	Y	Y
Code.org, (2021) AI and ML Module	Y	N	N	Y
Udacity, (2022) AI for Beginners	N	N	Y	Y
Coursera, (2022) Introduction to AI	N	N	Y	Y

Table 19: Introductory AI Courses Comparison Table

The main features for comparison within Table 19 were determined by identifying the key features within MetaLearning tutorials to identify the extent to which these were offered within other educational tools. As displayed in Table 19, MetaLearning is the only tool to support all of the categories. Specific mathematics content pertaining to AI was also offered within the AI and Machine Learning module, however this content was basic and only pertained to the calculation of accuracy. The courses offered by Udacity (2022) and Coursera (2022) both had pre-requisites for a level of mathematical understanding of calculus, probability and linear algebra.

All of the tools presented information using a variety of methodologies. However, study skills support was lacking from all of the courses apart from MetaLearning, this was an issue outlined in Section 5.2 detailing the lack of pedagogical consideration, leading to high dropout rates, particularly with MOOCs.

6.2.1 Evaluation of Numbas

As discussed in Section 5.4 (p.141), Numbas (Newcastle University, 2015c) was chosen as the host for the MetaLearning tutorials over existing systems, including prevalent VLE's such as Canvas (Instructure, 2022) and BlackBoard (BlackBoard Inc., 2022). The decision to use this infrastructure enabled ease of access for research participants as a direct link to the tutorials

could be provided and easily accessed. As well as integration within the VLE used by the participating universities. The numerical functionality within Numbas enabled ease of inclusion of mathematical notation, ensuring the creation of the Maths and Stats tutorial was relatively uncomplicated. Numbas has also undergone extensive usability testing and as discussed in Section 5.5, has been evaluated in a range of educational settings.

One potential limitation of the use of Numbas as the host for the MetaLearning tutorials is the lack of data collected pertaining to learning analytics. This data can be obtained if the tutorials are incorporated within a VLE, however, no additional data is collected when used in standalone mode which is how the majority of users participating in this evaluation accessed the tool. This method of accessing MetaLearning complies with requirement R.7, outlined in section 5.3.2. The lack of learning analytics is not deemed an issue due to the extensive evaluation of MetaLearning with a variety of users discussed within the following sections.

6.3 Recruitment of Participants to Review MetaLearning

To determine participants to evaluate MetaLearning, lecturers who consented to take part in the case study were approached first to ascertain if they were willing to trial the tool with their students. Lecturers who were identified as teaching this domain within the first phase of this research (as discussed in Chapter 3, p.61) were also contacted regarding MetaLearning. As well as asking the lecturers if they wished to introduce the online tool to their students, the lecturers were also approached regarding an interview to determine their view on the perceived usefulness of the system. Professionals within the computing domain who were not specialists in AI were also contacted to gain a range of opinions. Although these participants would not have the expert domain knowledge, they would have a wealth of experience teaching difficult computing concepts at a range of educational levels. Overall, three participants were recruited, two were from university A and one from university B of the case study institutions (Chapter 3, section 3.3, p.66).

6.4 Student Questionnaire

6.4.1 Student Questionnaire Methodology

Student participants who trialled the tool were from two of the universities from the case studies, university A and university B. However, the tool was trialled with a different cohort of

students and a different lecturer at university A. This was a senior lecturer in Data Science and Machine Learning. Within university A, two different sets of students used the tool. Undergraduate students received a guest lecture on Machine Learning, as part of their professional development. The tool was used as a blended learning approach for the Machine Learning session where the students were instructed on the basics of Machine Learning and were then required to complete the Overview of Machine Learning and Maths and Stats for Machine Learning tutorials. After completion of the tutorials, the students were able to ask any questions they had based on the material they had covered within the online tool.

Postgraduate students at university A also used the tool within their Deep Learning module to help them boost their understanding of particular topics, particularly mathematics and statistics. Links to the tutorials were provided as additional resources within the lecture slides. Both sets of students were notified of an optional questionnaire to provide feedback on their experience using MetaLearning.

Students enrolled on the undergraduate Artificial Intelligence module at university B were offered MetaLearning as an additional resource and were given time during their practical sessions to complete the tutorials. Due to the variation of content covered on the AI module, students were encouraged to complete all of the tutorials.

The student questionnaire was online and hosted through Newcastle University (2021) Form Builder and consisted of nineteen questions (Appendix J). Two of the questions centred around the usefulness of the tool in furthering understanding. Participants were asked if they felt the system had helped with their understanding of AI and Machine Learning and on a scale of 1 to 5, with 1 being no improvement and 5 being greatly improved how much they feel the tool had improved their confidence in this domain. This would help identify if the tutorials and the strategies had any impact on their belief in their ability. The majority of the questions were either multiple choice or checkboxes. However, open text boxes were included when greater depth of response was required, for example with the question “Any specific content you would like to see included?”

One of the main focuses of MetaLearning was to improve mathematics and statistics educational provision but to also boost student confidence in these skills, as discussed in

Section 2.8.1, mathematics anxiety can cause issues with attainment. Therefore, mitigating against this issue was a key aim of MetaLearning. Three questions relating to this aspect of MetaLearning were included within the questionnaire. Respondents were asked to rate their confidence in their mathematics skills on a scale from 1 to 5, with 5 indicating they were very confident. Participants were also asked if they felt the Maths and Stats tutorial had improved their knowledge within this area. A knowledge survey (Appendix K) was included within the Maths and Stats tutorial to enable students to reflect and assess their confidence within specific aspects of mathematics pertinent to AI. Questions within the knowledge survey related to specific mathematical content that was included within the tutorial including “What is a matrix?” and “What is the main advantage of Bayes theorem?” It was important to gain feedback from the students on whether they found the knowledge surveys useful and whether the tutorial progressed their mathematics skills. Therefore, questionnaire participants were asked to rate on a scale of 1 to 5 (5= very useful) how useful they felt the knowledge survey was in helping them reflect on their mathematics and statistics knowledge.

The following question was designed to understand which topics within MetaLearning the participants found the most difficult. The questionnaire respondents were presented with a list of all the tutorials and were asked to select the one they found most difficult. The two subsequent questions related to strategies to aid users in their learning. The use of low stakes quizzing through incorporation of a question at the end of each tutorial page was utilised as a method for students to receive feedback and to boost their self-efficacy. Participants were asked within the questionnaire to rank how useful they found the inclusion of the low stakes questioning as shown in Figure 32 to discern whether this was an effective strategy.

On a scale of 1-5, how useful were the questions at the end of each tutorial? (5= very useful)				
1	2	3	4	5
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 32: Question 7 of MetaLearning review questionnaire

A question relating to the inclusion of visualisation within the tutorials was included to determine if the user felt this had aided their understanding of the topics on which this method was applied. The following subset of questions related to usability and overall

experience with the tool, these included questions related to how easy the system was to use and how likely they would be to reuse it as shown in Figure 33.

<p>On a scale of 1-5, how easy was the system to use? (1= very easy, 5= very hard)</p>				
1	2	3	4	5
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
<p>On a scale of 1-5, how likely are you to reuse the online tool? (1= I will not reuse it, 5= very likely to reuse it)</p>				
1	2	3	4	5
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 33: Usability questions from the MetaLearning questionnaire

The responses indicating how likely they would be to reuse it may potentially reveal how useful they found the tutorials as they may be reluctant to attempt the tutorials again if they found they weren't worthwhile. Questionnaire respondents were also provided with a text box so that they could list any features they felt could be improved upon as shown in Figure 34. Participants were also asked if there was any additional content needed, this was to give insight into the aspects of AI that students were interested in learning but that were not yet present in the tool.

Any features in the online tool which you feel could be improved upon?

Figure 34: Question on features which can be improved on

One of the key themes of this research was discovering if students have mathematics anxiety or low confidence within their technical skills. As evidenced in Chapter 4, students in the case studies found the practical aspects of their modules more difficult than the theory. A number of the tutorials contain further resources including Python implementation or code snippets

to assist the users in linking the theoretical concepts with practical implementation. Within the questionnaire, participants were asked to rate their confidence in applying their knowledge within a practical context, as shown in Figure 35, this would give an indication of student feelings on their level of practical skills. Participants were also asked to rate their motivation relating to continuation of their studies within the AI domain.

On a scale of 1-5, how confident do you feel in applying your knowledge in machine learning/AI within a practical context? (5= very confident)				
1	2	3	4	5
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

Figure 35: Question relating to confidence in applying knowledge in a practical context

To conclude the questionnaire, information pertaining to respondent demographics was requested to determine representation within the academic pipeline (as discussed in Section 2.4), like all questions in the questionnaire, responding was optional. Students were asked which level of study (undergraduate/postgraduate/other) they were currently undertaking and which year of study they were in. As the tutorials serve as an introductory course for AI and its sub-domains, the material is suitable for a variety of educational levels. However, it is important to distinguish any difficulties or differences students may have relating to their educational level as this will help towards the framework of best practice and specific learning strategies.

6.4.2 Student Questionnaire Analysis Methods

The online questionnaires completed by students were transferred into SPSS (IBM Corp, 2019) from the Newcastle University Form Builder (Newcastle University, 2021). As the majority of the questionnaire data was categorised as either nominal or ordinal, non-parametric statistical methods were used for analysis. Non-parametric tests such as the Mann-Whitney test use the rank order of observations rather than requiring the data to follow a particular distribution as is the case with parametric tests (Altman and Bland, 2009).

Descriptive statistical techniques were used such as frequency analysis and crosstabulations (refer to Section 3.4.1, p.74). Frequency analysis was employed on the majority of the data to

determine the most frequent responses to the questions. Crosstabulation was used to determine the intersection of responses, for example to understand if there was a difference in mathematics confidence between genders or which tutorial they thought was the most difficult. This method was also used to look at the differences between the educational levels and responses to particular questions, enabling identification of any variation in difficulties students were encountering. These findings will contribute to the fulfilment of RO1 (p.6) in identifying the barriers students are facing when undertaking education within this domain. As well as assessing how and to what extent MetaLearning has assisted students in their learning (RO2.b, p.6).

Spearman rank order correlation (see Section 3.4.1, p.76) for ordinal data was used to examine if there was a correlation between how useful the respondent found the knowledge survey and their confidence level in their mathematics knowledge. This technique was also used to determine if there was a correlation between how much (if at all) the tool had improved the respondents' confidence within the AI domain and how motivated they were to continue studying this topic. These findings would inform RO2.c (p.6) in determining if the mitigation strategies included within MetaLearning improved the students metacognition and self-efficacy. Finally, this type of correlation was used to determine if there was any relationship between how easy the student found the tool to use and whether they stated they would use it again.

In respect to inferential statistics, the chi-square test for independence (see Section 3.4.1, p.78) was used to explore whether there was an association between a number of categorical variables. These included the following items: (i) respondent gender and the tutorial they found the most difficult, (ii) the study level of the participants and the tutorial they found the most difficult and (iii) if they felt the system had helped with their understanding of AI and Machine Learning. These findings would provide further insight into the barriers students are facing (RO1), the threshold concepts due to identification of the most difficult tutorials (RO2.a) and how successful the resource is in assisting learning (RO2.b).

6.4.3 Student Questionnaire Results

Overall, there were 77 responses to the student questionnaire. The respondents were from two different institutions, university A and B and were at a range of educational levels. Within

university A, first year Computer Science undergraduates used MetaLearning as part of a blended learning class introducing them to Machine Learning. The tool was also offered as a resource to second year undergraduate students who were interested in completing a final year project within this domain. MetaLearning was also offered as an additional resource to postgraduate students completing a module in Deep Learning. Within the lecture slides, links to the tutorials were available to compliment the material and offer the students additional material to bolster their learning. University B, a post-1992 university from the case studies (outlined in Chapter 3 and Chapter 4) trialled the tool as part of the Artificial Intelligence third year undergraduate module. Table 20 outlines the overall educational level demographics for the respondents. The majority of students (at 84.4%) were studying an undergraduate degree, further analysis indicated that 76.6% of these were in their first year of studies.

What level of study are you currently undertaking?

		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	1 UG	65	84.4	84.4	84.4
	2 PG	11	14.3	14.3	98.7
	4 No answer	1	1.3	1.3	100.0
	Total	77	100.0	100.0	

Table 20: Participants level of study UG (Undergraduate), PG (Postgraduate)

Information pertaining to participants' age and gender was collected. The majority of students at 87% (67) were in the age category 18-22, 7.8% (6) were 23-27 and 3.9% (3) were between 28-32 years old. 1.3% (1) of respondents preferred not to say. As figure 36 shows, the majority of respondents were male at 66.23% (51). Although the number of self-identified female respondents is low at 27.27% (21), only 20% of AI professionals in the UK are female (Young, Wajcman and Sprejer, 2021). Data relating specifically to the gender of students studying AI is currently unavailable. However, the gender split for Computing programmes, on which AI is a key topic is 20% (HESA, 2021) so these figures are in line with what was expected.

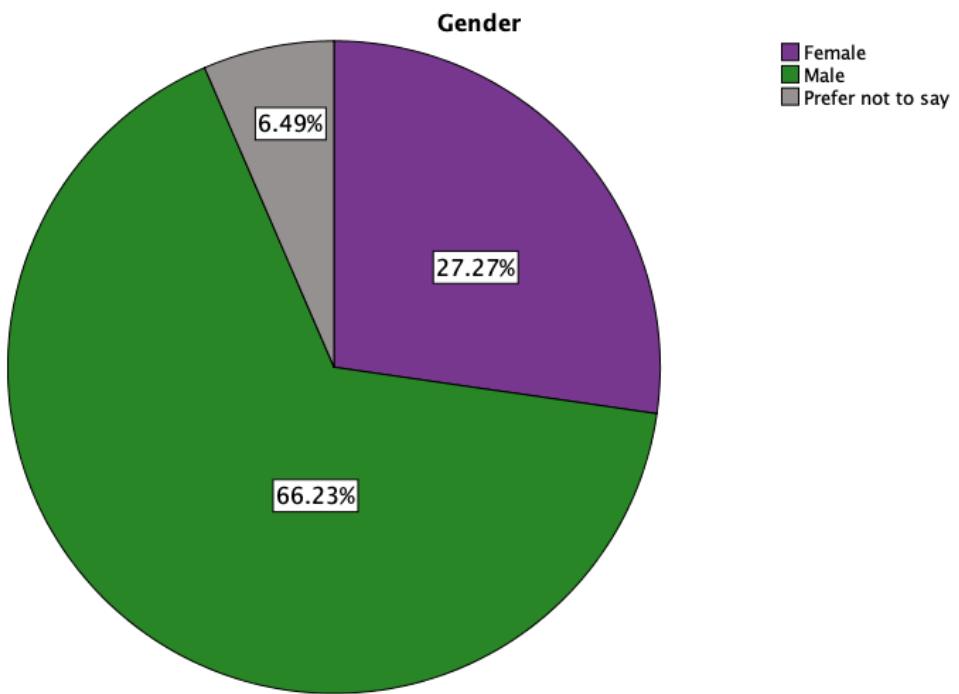


Figure 36: Student gender demographics from the student MetaLearning questionnaire

The respondents were asked whether they felt that MetaLearning had helped them with their understanding of Machine Learning and AI. 92.2% (71) chose yes, 6.5% (5) said that they weren't sure and 1.3% (1) did not answer this question. None of the participants that responded said that it hadn't helped with their understanding. Fisher's exact test (Sprent, 2014) was used to determine if there was an association between the level of study of the participants and how useful they found the tool. It was hoped this would provide information on whether the undergraduate or postgraduate students found the tool the most helpful. However, after running the statistical test, no clear association was found between level of study and perception of the tool's helpfulness.

Students were also asked on a scale of 1-5 (with 5 being greatly improved and 1 being not improved at all) how much they felt the tool has improved their confidence in this domain.

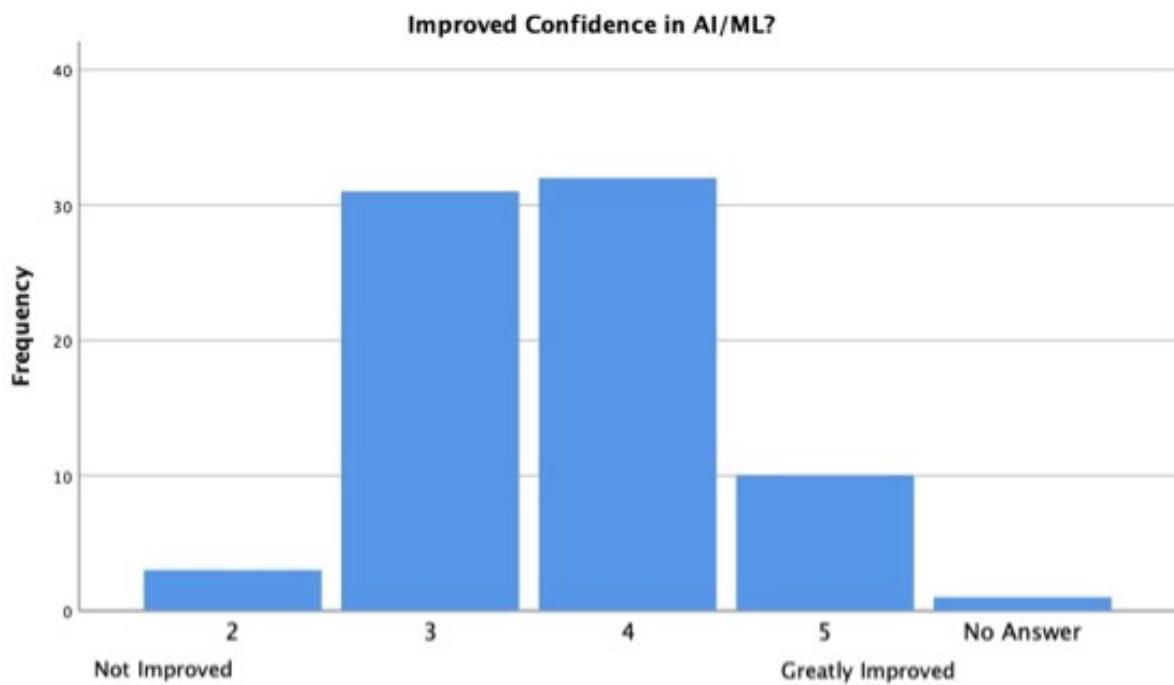


Figure 37: Student views on whether MetaLearning improved confidence

Figure 37 shows the responses to question 2 – how much do you feel the tool has improved your confidence in AI and ML (Machine Learning)? Out of the 77 respondents, 32 (41.6%) rated their improvement in their confidence as a 4, 31 respondents (40.3%) rated their improvement at a 3. The majority of the participants expressed that MetaLearning had led to some improvement in their confidence relating to AI and Machine Learning. In relation to confidence, the students were also asked to rate their mathematics confidence on a scale of 1-5, with 5 being equivalent to very confident.

As shown in Figure 38, the respondents on average had quite a high level of confidence in their mathematics skills. The most frequent score on the scale was 4 at 45.5% (35). However, 9 (11.69%) of the participants rated their confidence below a 3 which indicates that there potentially might be some students with mathematics anxiety.

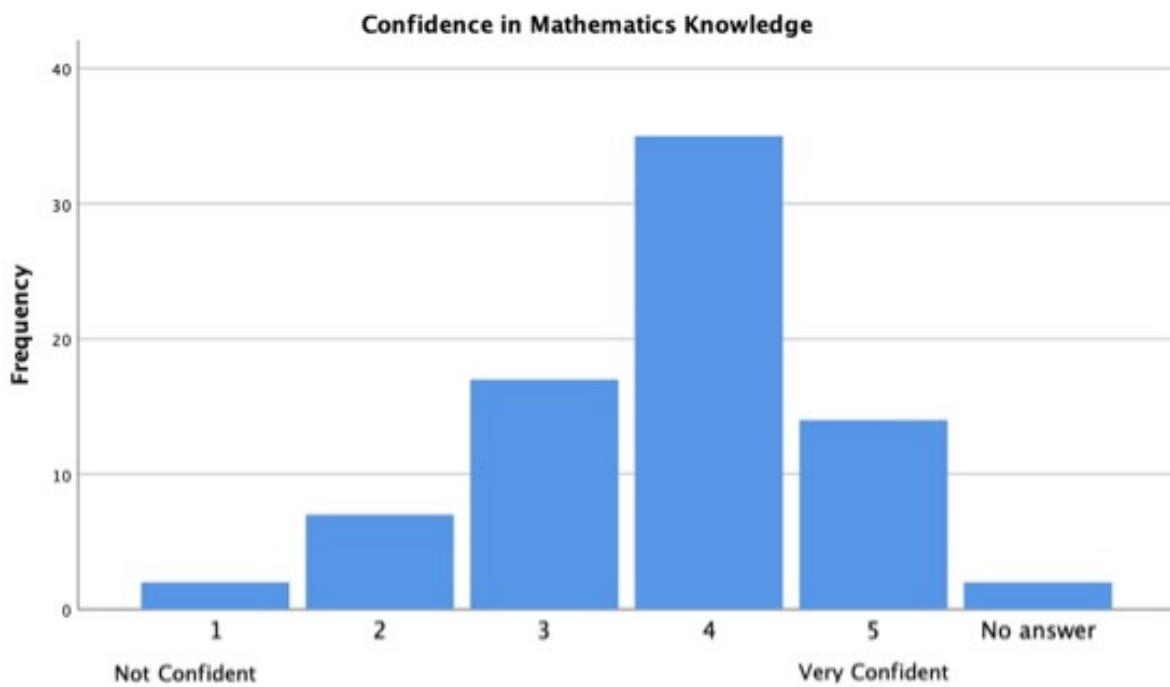


Figure 38: Responses relating to confidence in mathematics

Determining how informative the users of the tool found the Maths and Stats tutorial was important to determine its effectiveness and areas for improvement. The majority of participants, as shown in Table 21, felt that this tutorial had improved their knowledge in this area. However, 31.2% stated that they were not sure whether it had impacted their understanding or not, this uncertainty on their knowledge may be a reflection of their metacognitive skill level and self-efficacy in their mathematics skills.

Maths and Stats tutorial impact on knowledge

		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	1 Yes	41	53.2	53.2	53.2
	2 No	10	13.0	13.0	66.2
	3 Not sure	24	31.2	31.2	97.4
	4 No answer	2	2.6	2.6	100.0
	Total	77	100.0	100.0	

Table 21: Responses relating to improvement in Maths and Stats

The final question pertaining specifically to the students mathematics understanding was based on whether they found the knowledge survey useful as a reflection tool. Similar to the previous questions, this was scored on a scale from 1-5 with 5 indicating they found it very useful and 1 indicating they did not find it useful at all. Figure 39 shows the frequency response

to this question, indicating that the majority of students did find the knowledge survey assisted them in reflecting on their mathematics and statistics knowledge.

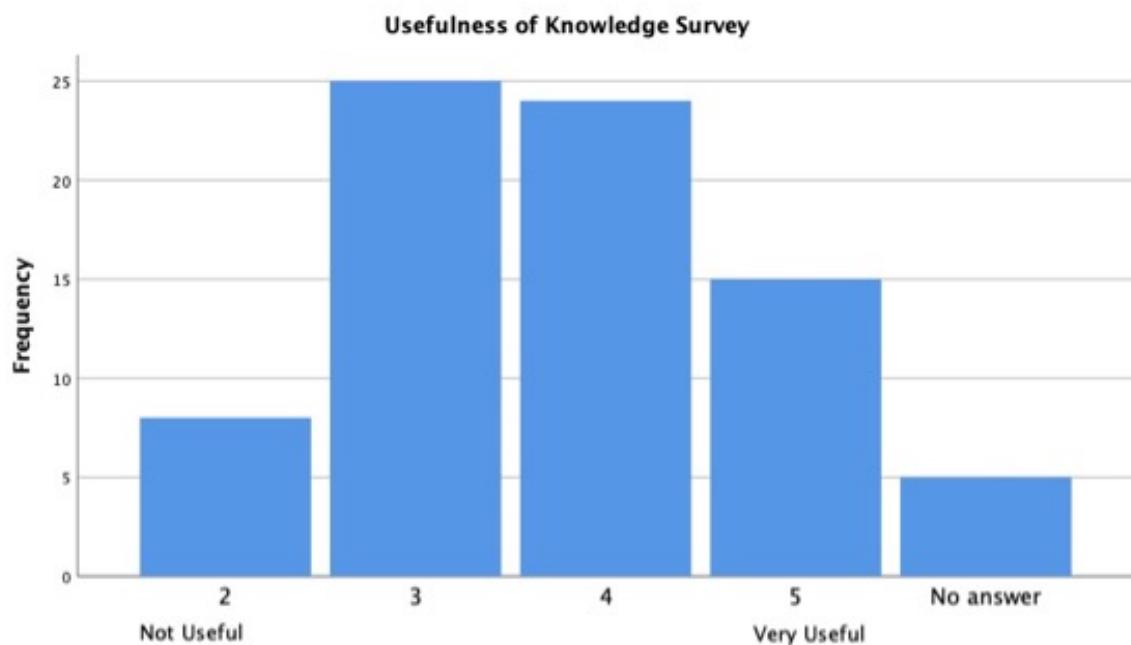


Figure 39: Frequency response to knowledge survey question

Spearman's Rank-Order Correlation (Zar, 2005) was applied to the variables relating to the knowledge survey and mathematics confidence to determine if there was any correlation between how useful the participant found the knowledge survey and how they rated their confidence in mathematics. A medium strength positive correlation was found, with the coefficient (ρ) equal to 0.389. The coefficient of determination was calculated, this shows that perceived usefulness of the knowledge survey helps to explain 15% of the variance in respondent's confidence in their mathematics.

Figure 40 displays the most frequent responses to the question asking respondents to choose the tutorial they found the most difficult. The most frequent response was the Maths and Stats tutorial at 33.8% (26). The next most frequent response was Machine Learning Algorithms with 16.9 % (13), closely followed by the Overview of Deep Learning tutorial at 15.6% (12).

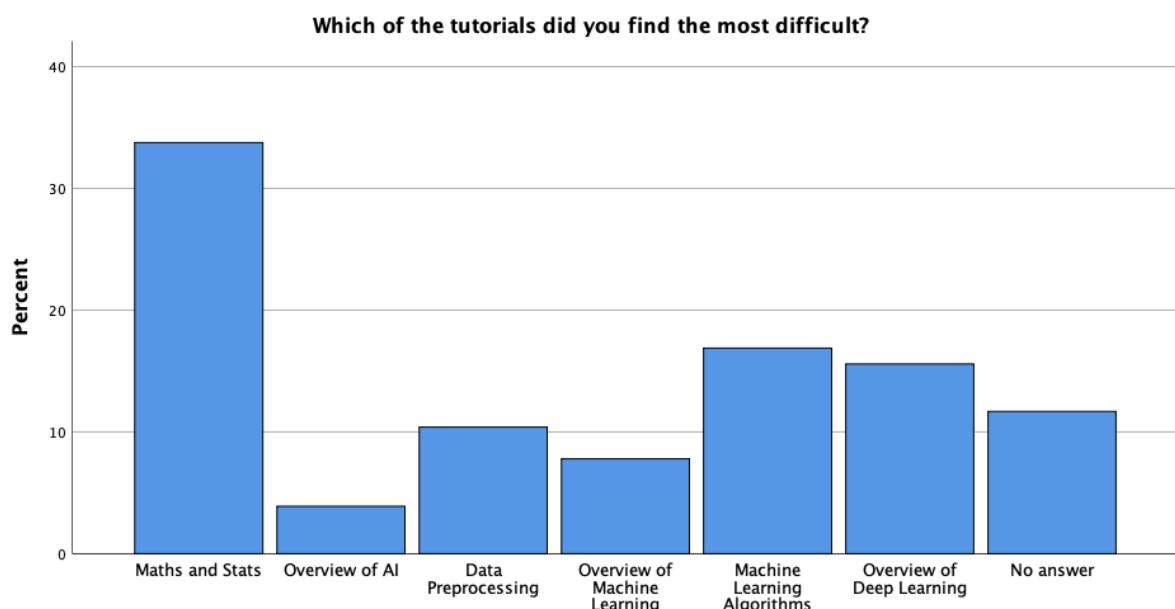


Figure 40: Percentage response to the Q: Which of the tutorials did you find the most difficult?

The chi-square test for independence (refer to Section 3.4.1, p.78) was applied on the gender and difficult tutorials variables to determine if there was a relationship between the student's gender and which tutorial they found the most difficult. However, the data violated the assumptions of this test due to the expected frequency in each cell being too low. Therefore, Fisher's Exact Probability Test was applied instead, which is applicable for smaller sample sizes. Fisher's Exact test works based on a null hypothesis that there is no association present (Freeman and Campbell, 2007). This test returned a value of 6.036 resulting in a p value of .953 indicating that there is not a significant relationship between the student's gender and which tutorial they found the most difficult. Fisher's Exact Probability Test was also used to determine if there was a relationship between the student's study level (e.g. undergraduate/postgraduate) and which tutorial they found the most difficult. This test returned a value of 14.291 resulting in a p value of .273, indicating that there was not a significant relationship.

Low stakes quizzing, in the form of questions at the end of each individual tutorial page was employed as a mitigation strategy to help improve student confidence and comprehension of their knowledge level within this particular topic. To determine how successful this strategy was, the students were asked to rate the questions usefulness on a scale of 1-5 (1 = not useful, 5= very useful). Figure 41 shows the percentage responses, indicating that the majority of students found this strategy useful as the highest response was '4' at 39% (30 respondents).

To assist the students in their learning of some of the complex topics within this domain, visualisation of difficult concepts was embedded within the tutorials. To determine how useful the users found this, they were asked within the questionnaire whether they felt visualisation had aided their understanding. 69 out of the 77 (90%) respondents stated that yes they did think visualisation had aided their understanding. This finding correlates with current educational best practice within this field outlined in Section 2.6.2 (p.43). Six of the respondents replied that they were not sure whether it had helped and 2 respondents did not answer the question.

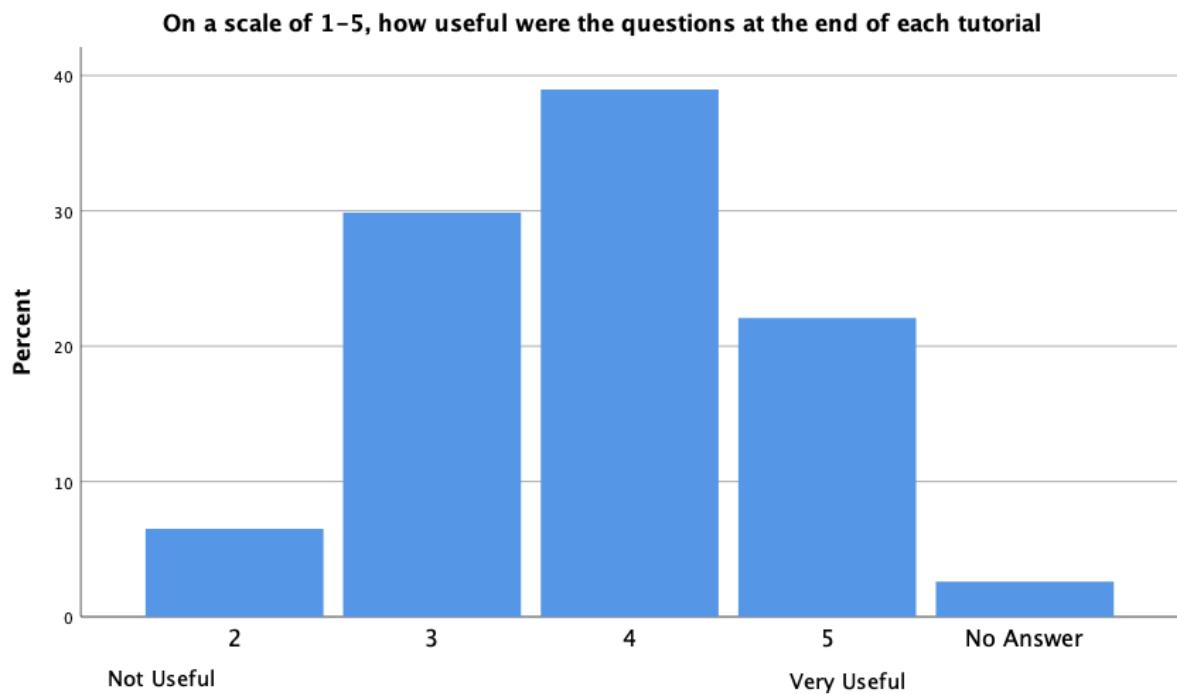


Figure 41: Responses to the questions relating to the low-stakes quizzing strategy

Questions pertaining to usability were important to determine how easy MetaLearning was to use and ultimately whether the users had found enough value in the tool to revisit it. The results from the question 'on a scale of 1-5 how easy was the tool to use' indicated room for improvement in the usability of MetaLearning due to the variation in responses to this question.

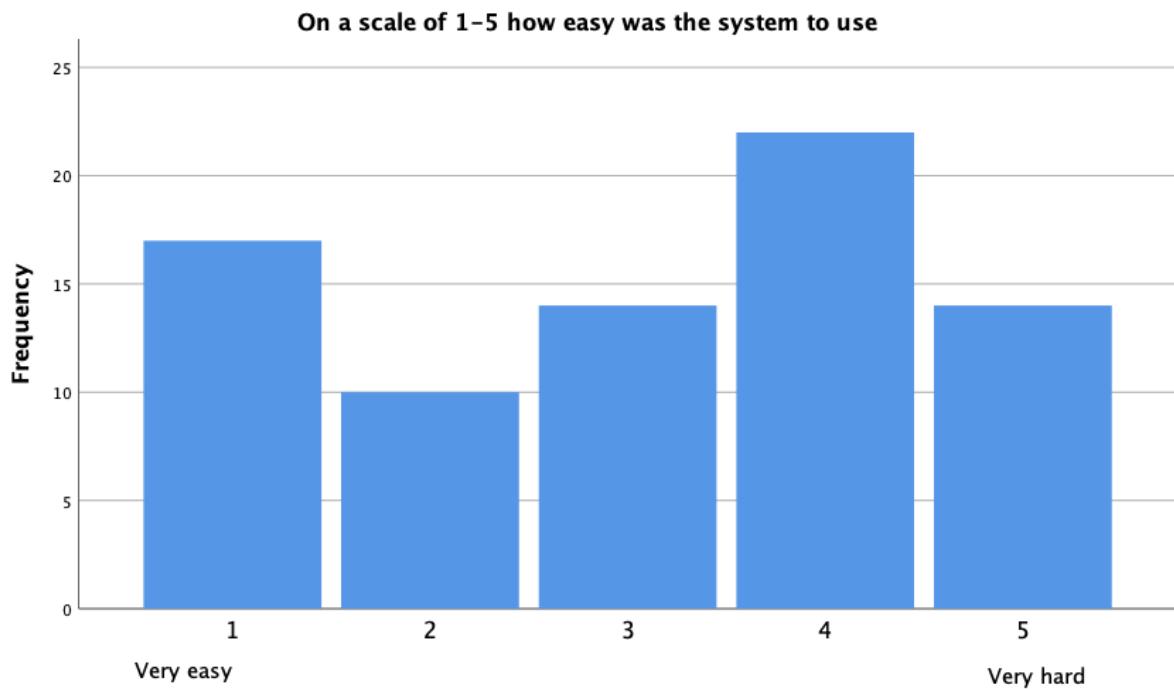


Figure 42: Frequency response to the question 'On a scale of 1-5 how easy was the system to use?'

As results show in Figure 42, the most frequent response was a '4' with 28.6% (22 respondents), indicating that participants found MetaLearning quite difficult to use. However, the next most frequent response at 22.1% (17) was a '1', which signifies these users found the system very easy to use. Therefore, there was some disparity amongst users. One possible explanation for this variation is that questionnaire participants misunderstood the rating scale. However, due to the relatively high responses within the middle of the scale, this does indicate that usability of the system should be improved. The results from the question asking participants how likely they were to reuse MetaLearning indicated that the majority of users would be likely to revisit the tutorials. The most frequent response on the scale from 1-5 (1=will not reuse it, 5= very likely to reuse it) was a '4' with 33.8% (26), indicating that participants were relatively likely to use MetaLearning again. Table 22 displays the frequency response to all points on the scale.

Spearman Rank Order Correlation (refer to Section 3.4.1, p.76) was used to determine whether there was a correlation between how easy the users found the tool and whether they were likely to revisit it, however no statistical significant correlation was found.

On a scale of 1-5 how likely are you to reuse the tool

		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	1 Not reuse	2	2.6	2.6	2.6
	2	7	9.1	9.1	11.7
	3	23	29.9	29.9	41.6
	4	26	33.8	33.8	75.3
	5 Likely to reuse	18	23.4	23.4	98.7
	6 No answer	1	1.3	1.3	100.0
	Total	77	100.0	100.0	

Table 22: Results from the question: how likely are you to reuse the tool?

There were two open-ended questions included within the questionnaire to allow respondents to advise upon any features within the system which could be improved upon and any specific content they would like to see included. The response rate for both of these questions was low, with 89.6% of students not answering the question relating to feature improvements and 94.8% of respondents choosing not to list any content they would like to see included.

Responses to the question relating to features which could be improved upon included “greater opportunity for interaction” and the “ability to zoom in on content”. However, someone commented that MetaLearning “has a user-friendly interface.” This variation in response relating to the user interface mirrors the findings on how easy participants found the tool to use, discussed on p.187. Therefore, further development of the user interface of MetaLearning will be required for future iterations. The responses relating to which content they would like to see added included “more mathematics examples”, “additional code examples” and to “further define the differences between the different types of machine learning”. These suggestions will be considered for inclusion within the next iteration of the tutorial set.

Results from the case studies outlined earlier in Chapter 4 (Section 4.5.4, p.120) indicated that students had difficulty with the practical aspects of this domain, therefore respondents were asked how confident they felt in applying their knowledge within a practical context. Respondents were asked to rate their confidence level on a scale from 1-5, 1 indicates that they are not confident, 5 denotes they are very confident. Figure 43 shows the variation of confidence levels, the most frequent response was ‘3’ with 33.8% (26), however 28.6% (22) of

respondents indicated their confidence level was below a '3' which indicates that some of the students have a lack of confidence with their practical skills.



Figure 43: Frequency response indicating how confident respondents are applying their knowledge in a practical context

Participants were also asked how motivated they were to continue studying AI and Machine Learning. On a scale of 1-5, with 5 indicating "very motivated" and 1 indicating "not motivated" over 66.3% (51) of respondents responded with either a '4' indicating they were quite motivated or a '5' showing that the students were very motivated to continue studying this domain. Spearman Rank Order Correlation (discussed in Chapter 3, Section 3.4.1, p.76) was employed to determine if there was a correlation between how much the participant felt MetaLearning improved their confidence in this topic and how motivated they were to continue studying this domain. The test showed a positive correlation, $r = .512$, with a coefficient of determination of 26.2%, indicating that improvement in confidence helps explain 26% of the variance in respondents scores on how motivated they were to continue their studies. This finding relates to the work of Lee (2009) as discussed in Section 2.8.1 (p.49) who discussed the correlation between confidence, self-efficacy and educational motivation.

Respondents were given the option to add any additional comments relating to MetaLearning and their experiences using it. 94.8% of questionnaire respondents chose not to answer this

question. However, 4 out of the 77 (5.19%) respondents did, 3 of the comments were positive towards the online tool and 1 was a negative response. The comments are listed below:

“Really enjoyed learning from your resources. This has reinforced my interest in AI and machine learning as this is why I chose to do computing at university”

“The tutorials seemed very good and all aspects were well explained”

“You taught me matrices and maths better than my GCSE maths teacher (I never did A Level Maths) so learning about matrices was very new and interesting as well as the refresher on algebra”

“Please remove the Maths and Stats for Machine Learning it is completely unnecessary. The focus is AI not A Level Maths”

The last comment indicates that this student has not understood the context for the inclusion of the mathematics and statistics tutorial and has not made the connection between the mathematical topics and its relevance to understanding the theory of Machine Learning. This finding indicates that the connection between the mathematics and statistics content and its relation to the Machine Learning content needs to be more explicit to ensure the users fully understand the relationship and foundational role mathematics plays within this domain.

6.4.4 Maths and Stats Knowledge Survey Results

The Maths and Stats knowledge survey (in Appendix K) covered most of the content included within the corresponding tutorial within MetaLearning. Students were asked to note their confidence level in answering the survey questions accurately. The confidence scale had 4 points (1 = Very confident, 2 = Somewhat confident, 3 = Not sure and 4 = Not at all confident). Overall, there were twenty-one questions pertaining to the content covered in the Maths and Stats tutorial. Example questions included “What is a matrix?”, “Explain what a partial derivative is” and “Describe the different types of regression.” The full set of questions are shown in Table 23.

The knowledge survey consisted of an anonymous online form, which users of MetaLearning were prompted to complete before and after undertaking the tutorial. One of the unintended negative consequences of fully anonymising the knowledge survey was that it was not possible to distinguish between the entries which were completed by the same participant, therefore indicating which form was completed before or after the tutorial. This unfortunately means that the efficacy of the Maths and Stats tutorial cannot be determined from this form of data collection, however Table 21 indicates that this tutorial was effective in improving users' knowledge and Figure 38 indicates that the majority of students, 45.5%, self-reported a confidence level of 4 in mathematics and statistics. This is a relatively high confidence rating as 5 on the scale denoted the student was very confident in their mathematics and statistics knowledge. The results from the questions pertaining to mathematics confidence in the student questionnaire will be compared with the findings from the knowledge survey to determine whether students are able to accurately reflect on their perceived level of confidence and whether the results tally.

Overall, there were 18 entries for the Maths and Stats knowledge survey. Frequency analysis was used to determine which of the questions the participants felt the least and most confident in answering. The questions which the most participants rated as "very confident" in answering were "What is a matrix?" and "Solve a linear equation." Both matrices and linear equations are on the GCSE curriculum (Department for Education, 2014) therefore participants with this level of mathematics attainment should have some form of prior knowledge of these topics.

Out of the twenty-one questions, nine of the questions (43%) got a majority confidence level of either "not sure" or "not at all confident." All of the nine questions either related to probability or statistics, Table 24 details the nine questions which the students had low confidence about answering and the percentage who answered with the response shown. Table 24 shows that out of this subset of questions, the ones relating to probability distributions were the ones in which the students reported the lowest confidence level in being able to answer.

Maths and Stats Knowledge Survey Questions
What is a matrix?
Transpose a matrix
Given an exponential and a logarithm, can you explain the relationship between the two?
How is the product of a number and multiplication related?
Explain the difference between linear and non-linear functions
Solve a linear equation
How does a derivative relate to change in a function?
Discuss the difference between the chain rule and the product rule for derivatives
Explain what a partial derivative is
Explain the difference between prior and posterior probability
Generate a short explanation of the relationship between probability and Machine Learning
Calculate the sum rule for a mutually exclusive event
What is the main advantage of Bayes theorem?
What is meant by Bayesian inference
Which probability distribution should you use if you have limited prior knowledge about what form a distribution should take?
Explain the relationship between the Bernoulli probability distribution and the Binomial distribution
Given a dataset, apply the Pearson Correlation Coefficient using Python
Explain why curve fitting is used
Describe the different types of regression
Explain how MSPE (Mean Squared Prediction Error) is related to regression
Given a dataset, calculate the standard deviation

Table 23: All questions from the knowledge survey

The question relating to applying the Pearson Correlation Coefficient in Python was included within the knowledge survey to give an indication of how confident participants were in applying their mathematical knowledge within a programmatical environment. The results show that the students had a lack of confidence in their ability to successfully execute this task as 44% of respondents stated they were not at all confident.

Question	Confidence Rating
Explain the difference between prior and posterior probability	Not sure (44%)
Calculate the sum rule for a mutually exclusive event	Not sure (44%)
What is the main advantage of Bayes theorem	Not sure (39%)
What is meant by Bayesian inference	Not sure (44%)
Which probability distribution should you use if you have limited prior knowledge about what form a distribution should take	Not at all confident (39%)
Explain the relationship between the Bernoulli probability distribution and the Binomial distribution	Not at all confident (44%)
Given a dataset, apply the Pearson Correlation Coefficient using Python	Not at all confident (44%)
Describe the different types of regression	Somewhat/Not sure (33%)
Explain how MSPE (Mean Squared Prediction Error) is related to regression	Not at all confident (33%)

Table 24: The questions which received the lowest overall confidence rating

6.4.5 Student Results Summary

Two universities reviewed MetaLearning, a Russell Group university (University A) and a post 1992 university (University B). Alongside the variation in universities, the cohorts who trialled the tool were a mix of both undergraduate and postgraduate students. Some of the students were currently completing a module within the AI domain and some students had not experienced any formal education within this domain yet. This variation in cohort is beneficial to understand how accessible and useful the tutorials are for students at varying educational stages and with differing educational backgrounds. As well as the effectiveness of the mitigation strategies (RQ3, RO2.c), as educational background has been shown to influence self-efficacy (as discussed in Section 2.8).

Out of all respondents, 92.2% said MetaLearning had helped with their understanding of Machine Learning. This finding demonstrates that use of MetaLearning had assisted the students in learning topics pertinent to this domain and fulfils RO2.b (p.6). Alongside assisting students with their learning, the respondents also felt that using MetaLearning had improved their confidence in their Machine Learning/AI knowledge. This improvement in confidence could have a positive impact on a student's metacognition as they will be able to more accurately reflect on their current level of understanding due to their new level of self-efficacy (RO2.c, p6).

One of the potential barriers identified as part of RO1 (p.6) was a lack of mathematics knowledge and potential mathematics anxiety. Findings from the students' questionnaires relating to confidence in mathematics indicated that the majority of students had a high level of confidence. However, 12% of respondents rated their confidence below a 3 on the 5-point rating scale (with 5 = very confident) identifying that confidence in mathematics was an issue for these students. The Maths and Stats tutorial may be a potential mitigation strategy, alongside the knowledge survey, to combat these confidence issues as 53.2% of respondents found that this tutorial had helped their understanding. The respondents also indicated that they found the knowledge survey useful in helping them to reflect on their current understanding of mathematics and statistics. The findings from the knowledge survey highlighted a particular issue and lack of confidence answering questions relating to probability, indicating that this may be a potential barrier (RO1).

The strategies included within MetaLearning as part of RO2.b and RO2.c to improve student's metacognition as well as assisting them when learning complex topics were both found to be useful for the students. The use of low stakes quizzing which was employed in all tutorials to challenge the users understanding of the material they had just learnt had a majority response of 4 (39%) on the 5-part scale (5 = very useful). The use of visualisation for complex topics was also deemed useful with 90% of respondents voting yes to the question asking if it aided their understanding. The positive response to these strategies demonstrates the effectiveness of these techniques in helping students learn this domain and the overall perception from the students that MetaLearning is a useful learning resource.

The following section outlines the lecturer view of MetaLearning, this will offer another perspective from the student view to determine if the lecturers have a similar view as students and whether they would recommend the tool to their students.

6.5 Lecturer Interviews

6.5.1 Lecturer Interview Methodology

Lecturers who took part in the case studies (Chapter 4, Section 4.5, p.101) and educators in Computing were contacted via email and asked to evaluate MetaLearning. The evaluation process involved evaluating two tutorials, Maths and Stats for Machine Learning and Machine

Learning Algorithms. Evaluation consisted of an online interview where the participant would be asked semi-structured questions to determine their experience using the tool. These tutorials were chosen for all of the interviews for consistency as they represented the main objectives of this research in assisting students with their learning of Machine Learning and inclusion of mitigation strategies to improve student metacognition and self-regulation.

An information sheet was included in the participation email detailing the purpose of the study, what participation would require and the type of information/data to be collected. Potential participants were given the opportunity to ask questions and enquire about any aspects of the study they were unsure of. Completion of the research information sheet/consent form was required before the interview could be conducted (a copy of this can be found in Appendix L).

The interviews were hosted online through both Microsoft Teams (Microsoft, 2021) and Zoom (Zoom Video Communications Inc, 2021) due to the current COVID 19 pandemic and restrictions in place. To aid the analysis process and to ensure accuracy of the transcription process, interview participants were asked if they consented to their interview being recorded. This was not compulsory and did not affect participation in the study if the participant declined.

Semi-structured interview questions were created for the interview to give some structure and to ensure information pertinent to the research objectives was obtained. Overall, there were thirteen questions in total (as shown in Figure 44). Questions relating to the participants current teaching situation were included to determine which educational level they taught and whether they taught some form of AI or Machine Learning. There were three questions relating to mathematics and the Maths and Stats tutorial, including whether they think this tutorial will help improve student's knowledge in this area and whether the knowledge survey will boost self-efficacy and confidence with their mathematics knowledge. Relating to the mitigation strategies included in MetaLearning, interviewees were asked whether they thought the low stakes questioning included at the end of each tutorial page would improve student confidence and help them better assess their knowledge level.

Interview Questions:

1. At what educational level do you teach?
2. Do you teach any form of AI/ML or Data Science?
3. Do you feel that the system would help your students in understanding Machine Learning?
4. What is your view of your student's mathematics and statistics knowledge?
5. Do you feel that the Maths and Stats for Machine Learning tutorial will help improve student knowledge in this area?
6. Do you think the Maths and Stats knowledge survey will be beneficial to improving student self-efficacy and confidence within this domain?
7. Do you think the use of low stakes questioning will improve student confidence in their understanding? And help students understand their knowledge level in the content?
8. Do you think the content is pitched at the right educational level for your students?
9. Any content which you would like to see included?
10. What do you think of the usability of the tool?
11. What features do you feel could be improved upon?
12. On a scale of 1-5 how likely (5= very likely) are you to recommend your tool to your students if they are interested in learning this area?
13. Anything else you would like to add?

Figure 44: Interview questions about MetaLearning

Two questions were included which related to the usability of the tool and any features the lecturers felt could be improved upon. Students were also asked this within their questionnaire. This question was included to determine if there was any disparity between staff and student perceptions of usability.

It was important to determine if the lecturers felt that the system would help their students to learn Machine Learning and how likely they would be to recommend MetaLearning to their students. Ultimately, the lecturers need to see some educational worth in the online tool for them to endorse its use for their students.

6.5.2 Lecturer Interview Analysis Methods

Upon completion of the lecturer interviews, the interviews were transcribed alongside any notes which were made in the interview. Thematic analysis (Section 3.4.2, p.79) was used to identify commonalities and patterns within the data pertinent to the research questions. Semi-structured interviews were completed with three participants.

6.5.3 Lecturer Interview Results

The three participants for the interviews were from two differing institutions one was a Russell Group university and the other university was ranked within the UK top 100 universities (The Guardian, 2020) and can be categorised as a post 1992 university. All three respondents had experience teaching some form of AI. Lecturer A from the Russell Group University mainly taught Deep Learning at postgraduate level and is also a prolific researcher within this domain. Lecturer B also from the Russell Group University had experience instructing on Machine Learning and Data Science to undergraduate students, however they had no prior experience being a module leader for this domain. Lecturer C from the UK top 100 university was currently teaching a postgraduate Data Science module. Although they were not the module leader their role was offering assistance to students and they were a newcomer to the field of Data Science.

All three respondents stated that they thought MetaLearning would help their students in understanding Machine Learning. Lecturer B said that they thought MetaLearning would also be useful for Computer Science students in general as it would serve as a refresher for the domain fundamentals. Lecturer C said that *“the descriptions are a good basic introduction if you’ve never done it before with complexity that, had you done it before, it would be a good refresher and reminder, a prompting.”*

All three lecturers said that their classes were comprised of students with varying levels of mathematical and statistical knowledge and the level depended on their educational background. Lecturer A noted that this variance in educational background meant that they had to include foundational mathematics within their lecture material to assist the students in understanding a particular element of Machine Learning or Deep Learning, therefore the online tool *“would be very helpful in a scenario like that, where the student is able to click a URL and go there and find out what they need to in terms of the foundational machine learning or maths and stats foundations that they need.”* Lecturers B and C also felt that the Maths and Stats tutorial would help improve their student’s knowledge within this area. Lecturer B stated that a lot of their students sign up for Machine Learning modules then want to drop out when they realise the extent to which mathematics is involved and that this tutorial would help ease them into the topic, especially the fact that they can work on it independently. This finding

highlights the importance of a student's mental model and preconceptions of the domain and the impact this can have on self-efficacy and motivation as discussed in Section 2.8 (p.47).

Lecturer views on the use of the knowledge survey to improve student self-efficacy and mathematics anxiety were mixed, Lecturer C was definitive in that they liked it and thought that it would “100%” help their students. Lecturer A thought that the survey may negatively affect student confidence, however they could see the benefit to it if the student completed it before and after the tutorial so that they could see how they have improved and what they had learnt, and that this confirmation and self-reflection would boost their confidence. Lecturer B thought that the benefit of completing the knowledge survey would depend on the student and their confidence level to start with, they thought this strategy of completing it before and after the tutorial would be most beneficial in conjunction with a teacher to then assist with their mathematics learning. These findings relating to the knowledge survey highlight the importance of the knowledge survey being completed as intended, therefore it may be beneficial to re-consider how this is incorporated within the tool to better signpost to users how to complete this task. It may also be worthwhile highlighting to lecturers who use the Maths and Stats tutorial that they may need to support students in this task.

All the lecturers expressed positive comments about the use of low stakes questioning to improve student confidence and understanding. Lecturer C advised that the students would like the fact that the questions do not have any impact on summative grades so they can “*just have a play around and figure out if something’s wrong without worrying about it impacting their final grade.*”

The interview participants also felt that the content was pitched at an appropriate level for their students. Lecturer A advised that “*the type of material that you have in your tutorials that I had a look at is excellent. It’s the right sort of things that I want my students to have a background knowledge of for getting into what I’m teaching.*” Lecturer C advised that they thought there was enough to cover the foundations without “*talking down to the students*” and that they liked that additional resources are provided so that students can read around the topic in greater depth.

Relating to the question on any specific content they would like to see included, lecturer A advised that more content relating to the metrics used to evaluate Machine Learning algorithms such as accuracy, precision and recall and the issues surrounding each of these would be useful. Lecturers B and C both mentioned the potential for a programming Machine Learning tutorial and how this might potentially assist the students in putting their mathematics and statistics knowledge within a practical context. However, lecturer C also mentioned that for students without prior knowledge in both mathematics and programming, the inclusion of this tutorial may be overwhelming and negatively affect their self-efficacy. All of the findings collected relating to additional content will be considered for future iterations of MetaLearning.

Relating to the usability of MetaLearning, all respondents agreed that there were no issues pertaining to its use. Lecturer B advised that they particularly appreciated the typeface used, *“as someone with dyslexia, I find it really easy to read.”* They also said that the formatting was user friendly in that it has whitespace and there is an appropriate amount of content on each page. All of the lecturers mentioned that there is potentially further scope for greater interactivity, for example lecturer A advised that it would be nice to break down the tutorials into smaller sections which are only revealed after the user has completed each section so that the material may be easier to digest and may hold learners with a shorter attention span.

Towards the end of the interview, the participants were asked to rate on a scale of 1-5 (5=very likely) how likely they were to recommend the tool to their students, Lecturers A and B both stated a 5 and Lecturer C advised a 3. However, Lecturer C said that if there was greater availability to pick and choose the content then they would give it a 4 or 5 as students are always looking for further material and practice questions on things they don’t understand. This issue will be addressed in subsequent versions of MetaLearning where a question bank will be compiled in which additional content will be available. The current methods of accessing the tutorials ensures that lecturers are in complete control and can choose which tutorials they share with the students.

To conclude the interview, the participants were asked if they had any additional comments they would like to make, Lecturer C declined. Lecturer A was complimentary and said that they thought *“this will be a very successful, high demand tool.”* Lecturer B reiterated their

earlier comment in that they think this tool will not only be useful for students studying AI but also for other Computer Science students as well.

6.6 UKICER 2021 Workshop

The UKICER 2021 conference was held online between the 2nd and 3rd of September. The workshop, 'Measuring the Difference Between Student and Staff Perception of Self-Efficacy and Confidence Using Online Tools' (Allen, Heels and Devlin, 2021) was undertaken alongside Laura Heels (ResearchGate, 2021), also from Newcastle University. The workshop was in part motivated by the shift in teaching as a consequence of the COVID-19 pandemic and how this highlighted and drew attention to inequalities of pedagogy (Peters *et al.*, 2020). These inequalities have exacerbated the already acknowledged issues relating to student's self-efficacy, confidence and metacognition (Yang *et al.*, 2021). The workshop was designed to examine best practice relating to online learning with discussion about its capacity to promote self-regulated learning and improve student confidence in their technical skills. To facilitate the discussion relating to strategies for improving self-efficacy and metacognition, the mitigation strategies implemented within MetaLearning were discussed. Participants were asked to review the perceived efficacy of these strategies as well as the online tool in general. Padlet (2021), a collaborative web platform, was used to host the activities and online discussions.

6.6.1 Running the Workshop

The workshop was hosted online through Discord (2021) which is a digital distribution and instant messaging platform. Within Discord, a chat channel was set up so that workshop participants could message, discuss and ask questions to other attendees and the workshop facilitators. There were eleven workshop attendees overall. The workshop lasted ninety minutes, therefore time was divided prior to the workshop to enable enough time for each of the five activities to be completed. To attempt to empower participants to freely provide their opinions and participate in discussions, it was important to set the tone for the workshop during the introduction. Participants were openly encouraged to share their thoughts either through the live stream or anonymously through the Padlet.

Interspersed between the activities were a number of discussions where the main topics related to that specific activity were introduced. These topics included self-regulation, metacognition and self-efficacy, concepts which have shown to help students organise their studies and become more independent learners. Amongst the barriers to learning and mitigation strategies to improve participation, inclusive language and the process of decolonising the curriculum was discussed. Guzdial (2021) advises that the field of Computing is not a meritocracy and that Computer Science education systems in particular are structured to “disadvantage students who are not like us and the students currently in CS.” As part of this discussion, a variety of resources were shared with participants such as the ACM’s Words Matter resource (ACM, 2021) and the Gender Decoder (Matfield, 2016).

The five activities were hosted on padlets which were pre-prepared with questions to facilitate the discussion, Table 25 outlines the exercises and the questions within the padlets. The first exercise was to determine what barriers the participants thought their students faced, this exercise was completed after an introduction to the concepts of self-efficacy, metacognition and self-regulation. The second exercise was designed to determine to what extent participants considered self-efficacy within their own practice and to share any good practice examples or less than beneficial experiences.

Once participants had considered barriers which they think their students faced and reflected on examples of good and bad practice they had encountered, the online tool for Machine Learning, MetaLearning was presented. Due to the specialisation of MetaLearning and the diverse expertise of the workshop attendees, the focus of the discussion surrounding the online tool was focused on the mitigation strategies, particularly the inclusion of the knowledge survey and low stakes questioning. Table 25 outlines the specific line of questioning. As well as the questions relating to the mitigation strategies included within MetaLearning, a vote was also included where participants could rate whether they thought the online tool would help with student understanding of Machine Learning.

Exercise Number	Exercise Name	Questions
Exercise 1	Barriers	<ul style="list-style-type: none"> - What barriers do you think your student is faced with? - What barriers do you think are overlooked? - What barriers are discussed too much?
Exercise 2	Self-Efficacy	<ul style="list-style-type: none"> - What are some best practice examples for consideration of student self-regulation, metacognition? - What are some “worst” practice examples for consideration of student self-regulation, metacognition and self-regulation? - How would you implement some of the discussed methods into your practice? - What resources would you need?
Exercise 3	Online Tool Review	<ul style="list-style-type: none"> - Do you think the tool would help with student understanding of Machine Learning? - What do you think we need to help students understand Machine Learning and do you think the tool covers it? - What do you think of the use of knowledge surveys as a method for improving self-efficacy? - What do you think of the use of low stakes questioning as a method for improving self-efficacy? - What are your thoughts on low stakes questioning when it comes to improving students understanding of a subject? - Additional comments?
Exercise 4	Inclusive Computer Science	<ul style="list-style-type: none"> - What do we need to do to make Computer Science more inclusive? - What do we need to be able to achieve this? - What do you think needs to be left alone? - Miscellaneous
Exercise 5	Evaluating the Online Tool for Inclusive Language	<ul style="list-style-type: none"> - Examples of good practice? - Areas that need to be improved to be more inclusive - Miscellaneous

Table 25: Workshop exercises and questions

Exercises 4 and 5 related to inclusive language and what we can do to make Computer Science learning materials more inclusive. Exercise 4 required participants to consider what we, as educators within Computer Science, can do to make the field more inclusive and to widen participation. Participants were then given the opportunity to build up experience in evaluating learning materials and resources for inclusive learning by evaluating the online tool for Machine Learning (MetaLearning) within a supportive environment. Due to the time limitations of the workshop, exercise five became more of an informal discussion on the use of inclusive language and how we can ingrain this within our practice.

To conclude the workshop, participants were asked if they had any further questions or points for discussion. Email addresses were made available to participants for any follow up questions. A follow up study is also proposed where the workshop will be run with students to ensure the views of both educators and students relating to the perceptions of self-efficacy are considered.

6.6.2 Findings

Following the workshop, the data from the padlets was downloaded for analysis. Due to the qualitative nature of the data, thematic and frequency analysis was completed to analyse the data (refer to Section 3.4.1, p.74). The organisation of the padlets prior to the workshop, including categorisation of the different discussion points enabled more efficient thematic classification. Figure 45 shows the completed padlet for exercise one.

Figure 45: Completed padlet for exercise one

The first exercise completed within the workshop pertained to barriers students face within their learning. Participants discussed a number of different barriers to student learning they had identified within their practice; four themes were identified through inter-rater reliability analysis with the other workshop coordinator. The first theme identified was in relation to threshold concepts within the Computing domain. The second theme related to the reluctance of students to ask questions in case they are deemed “*stupid questions*”. The participants noted that students perceive their question to be stupid “*in the sense that everyone else knows or understood already*.” This subject of concern relating to peer opinion also emerged within the third theme where participants identified that students comparing

themselves to each other can be a barrier, especially for students in a mixed ability class. One participant defined mixed ability as prior experience, not ability relating to grades or qualifications. Another participant added a discussion point suggesting that students assume that other students are high scoring because they “*are a genius or gifted.*” This can be a barrier to student self-efficacy and motivation if the students do not feel a high score is achievable or within their capability. The final theme was related to grades, and the expectation or perception that a particular grade is “*needed*”. This can lead to feeling of failure for the student if they do not reach the particular grade which they had targeted for.

Relating to the second discussion point within exercise one, “what barriers do you think are overlooked?” the main theme was around educational background and prior experience. One participant stated that “*beginners come in with so much different prior experience - not necessarily associated with formal education - that shape their skillset and ability to work with computational concepts.*” The idea that students have differing skillsets outside of formal educational concepts was also raised by another attendee who advised that students who are more introverted or shy can experience this as a barrier to learning due to this potentially preventing them from working with peers or accessing support from staff. Relating to prior experience, one participant raised the point that it is often assumed that students are conversant with different tools, particularly ones required for practical work.

Exercise two related to self-efficacy and sharing best practice and methods to improve student metacognition and self-efficacy to become more self-regulated learners. The main theme identified from this discussion was the importance of good quality feedback. One of the participants suggested a “*feedforward session...making targeted feedback points based on recurring errors across the class.*” Making students aware of the areas in which they need to improve whilst informing them that the errors they made were common within their cohort may not only improve their confidence but will also pinpoint to them specific areas in which they need to further their understanding. One participant also mentioned that a useful resource may be a list of key questions which can be used by students to reflect on their assignment. It was also raised that the classroom environment “*doesn't always work*” and that social politics can affect a student’s self-efficacy. Therefore, alternative modes of delivery may be more suitable for certain learning situations such as online and blended learning.

Exercise three related to two particular strategies for improving self-efficacy and metacognition which were included within MetaLearning. Participants were positive about the use of the knowledge survey, one participant said that it helped set the expectation of the student as they would know which questions they did not know the answers to in advance. However, one participant was less keen on the use of low stakes quizzing as they had experienced low participation in this strategy within their classes and they thought that it could cause stress for the weaker students and the students who are perfectionists. All of the questions contained within the tutorials are optional and anonymous if hosted as a standalone webpage, therefore students with lower confidence in their ability may wish to not partake in the quizzes or may take comfort from their answers being confidential. However, the option for restesting is a key feature of MetaLearning, therefore students have the option to build confidence in their knowledge through multiple attempts at answering the questions.

Findings specific to the use of MetaLearning as an educational resource for Machine Learning were positive. As Figure 46 shows, six of the participants voted with a thumbs up to the question “do you think the tool would help with student understanding of machine learning.” One participant also commented to say that the tool would help as a “*knowledge recap*.”

Do you think the tool would help with student understanding of machine learning?

Quick Vote

Click the like if you think the online tool would help student understanding of machine learning, click the dislike if you don't.

6 likes 0 dislikes

Anonymous 3d
Yes, as part of a knowledge recap.

Add comment

Figure 46: Results from the padlet vote from exercise three

Within exercise 3, participants were asked for their thoughts on what they think we need to help understand Machine Learning and whether they feel the online tool covers this. Again, responses were positive. One participant stated that *“You are making sure loads of prior knowledge is in place.”* Another participant advised that the tool is a good foundation for *“applied examples”*, demonstrating, for example how, the theoretical knowledge relates to the practical outcomes.

Widening participation and representation within Computer Science and particularly within the field of AI is imperative to ensure diversity within the teams developing these technologies and for democratisation of this domain. The purpose of exercise four was to gain a variety of perspectives from professionals within the field of Computing education on how we can make Computer Science more inclusive. Ideas included more prevalent role models and the idea of changing the “culture” surrounding the domain and raising awareness and responsibility that all students and staff need to take responsibility for their understanding and knowledge surrounding equality, diversity and inclusion.

6.7 Discussion

Determining both lecturer and student views on MetaLearning was imperative to understand how useful these particular users perceived the tool and to ascertain the perceived efficacy of the mitigation strategies. It was also important to establish whether student and lecturer views on the tool aligned and any aspects where there was a divergence in opinions.

The majority of students who trialled MetaLearning identified that it has helped them with their understanding of Machine Learning. This concurred with the lecturer view and participants of the workshop in which all lecturers interviewed felt that MetaLearning would help their students understand this domain. One lecturer also highlighted the potential usefulness of the online tool in general for Computing students. The majority of the student participants also noted an improvement in their confidence relating to their understanding of AI and Machine Learning.

An important aim of MetaLearning was to try and level up the disparities in students educational background and prior understanding relating to mathematics and statistics as

discussed in Section 2.6.2. This disparity was also noted by the lecturers reviewing the tool, within the UKICER workshop and also in the results pertaining to mathematics confidence in the student questionnaire. 53% of the students felt that the Maths and Stats tutorial had improved their knowledge within this area. Although the majority of students felt that they had made an improvement, 31% were unsure whether this tutorial had impacted their understanding or not, this uncertainty may indicate either a lower level of metacognitive skill in that they are not able to accurately reflect on their skill level or a potential issue with their self-efficacy in their mathematics skills. All of the lecturers who reviewed MetaLearning felt that the tutorial would help their students with their mathematics knowledge. In the interviews, Lecturer B also highlighted the lack of student understanding or mental model of what Machine Learning entails and that when students realise the extent of the mathematics and statistics involved, they tend to drop out of the module. However, Lecturer B advised that they felt the tutorial would ease their students into this domain and that it was particularly useful that the students could work through these tutorials independently.

The inclusion of the knowledge survey within the Maths and Stats tutorial was popular with the students reviewing the tool. The majority of respondents felt that it had helped them reflect on their understanding and knowledge level within this domain. The lecturers interviewed were cautious about the benefits of this strategy and felt that it would only be beneficial if it was used correctly or within a blended learning environment. However, the UKICER workshop participants liked this approach to improve self-efficacy and metacognition as they felt it helped set the expectations of the student as they would know prior to completing the tutorial which questions they did not know the answer to.

An interesting finding from the results of the completed knowledge surveys was that all of the questions where the majority chose the responses “Not sure” or “Not at all confident”, all of the questions related to either probability or statistics. This points towards a potential issue with the current education provision related to this specific domain and a possible obstacle to student comprehension and mastery of the Machine Learning domain.

The response to the use of low stakes quizzing as a mitigation strategy was positive overall. The majority of students voted a ‘4’ out of ‘5’ (61%) for the question ‘How useful were the questions at the end of the tutorial?’ All of the lecturers interviewed perceived the use of low

stakes quizzing as beneficial to their students. However, one of the participants in the workshop highlighted prior issues they had experienced with this technique relating to minimum participation and additional stress for students. This finding was particularly troublesome as it is the opposite consequence of its intended outcome. Although none of the students who reviewed this tool reported this issue, it will be important to monitor views of this strategy in future use.

Students found the Maths and Stats tutorial the most difficult out of all of the tutorials, this may be linked to the finding outlined by the lecturers that some students have a lack of prior knowledge pertaining to this domain. The students also found the Machine Learning Algorithms and Overview of Deep Learning tutorials difficult. Determining whether students found these tutorials difficult because they potentially contained threshold concepts or whether this was an issue inherent to these tutorials was important to determine whether content changes needed to be made. However, all of the lecturers interviewed who reviewed the tool said they thought the content was pitched at the correct educational level.

The scope for inclusion of further programming tutorials and code examples was raised by both students and lecturers. This is an area which will be considered for the next iteration of MetaLearning as the practical aspect of this domain have consistently been identified as an issue for students learning this topic both in Chapter 4 within the case studies (Section 4.5.4) and within the knowledge survey and student questionnaire. Within the student questionnaire participants were asked how confident they were applying their knowledge within a practical context, the majority answered as a '3' (on a scale of 1-5, with 5 = very confident), however there was variation in the responses with some students scoring themselves as low as '1' or '2'. Responses from the knowledge survey indicated that the majority of students (44%) who completed this were 'Not at all confident' in their ability to apply Pearson's correlation coefficient within Python. Two of the lecturers within their interviews mentioned the possibility of creating a specific programming tutorial as they felt this may help contextualise and embed the theoretical knowledge within a practical context. One of the workshop participants also highlighted the importance of "applied examples" for the purpose of assisting the students in understanding the correlation between the theoretical and practical knowledge. Therefore, further educational provision relating to the practical

application and execution of Machine Learning will be considered for inclusion into MetaLearning.

6.8 Summary

This chapter has outlined the review process undertaken for the online tool for Machine Learning, MetaLearning. To gain as wide a reach for this review process the tool was trialled with both students and lecturers and within a conference workshop. The overall view from both students and lecturers was that the online tool is helpful for students wanting to learn this domain and in particular students found the mitigation strategies of the knowledge survey and low stakes quizzes useful to their learning. In the next chapter we will revisit the initial research questions to determine if they have been addressed and also discuss the key findings from this research.

Chapter 7. Towards a Framework for Teaching Machine Learning

7.1 Introduction

This chapter reflects on the review of the current educational provision relating to AI and the potential barriers to learning identified through the case studies (Chapter 4, p.81). The results from the deployment (Chapter 5, p.130) and evaluation (Chapter 6, p.172) of the online tool for Machine Learning (MetaLearning) will also be discussed. The findings are discussed in relation to the research objectives and research questions outlined in Chapter 1 (p.5) and the literature review in Chapter 2 (p.9). These findings and analysis contribute to the framework outlined within this chapter.

The framework is focused around three core elements, the difficulties/barriers to learning this domain (RQ2, RO1, RO2.a, RO2.c), strategies to alleviate these difficulties (RQ3, RO2.b, RO2.c, RO2.d) and best practice for teaching this topic (RQ1, RO2.e). Each of these categories contain the key foci of this study including but not limited to educational background, mathematics anxiety, threshold concepts, online pedagogy and widening participation. Figure 47 provides a diagrammatic overview of the framework demonstrating how each of these contribute and feed into one another.

The chapter starts with an examination of the potential barriers within Machine Learning education identified through the analysis of the current education provision, outlined in Chapter 4, this addresses RO1, RO2.a, RO2.c and RQ2 (Chapter 1, p.5). This section of the chapter (Section 7.2) pertains to difficulties identified with mathematics anxiety and disparities in educational background as well as the first steps toward identifying the domain threshold concepts.

The next part of the chapter, Section 7.3, discusses the online learning tool for Machine Learning, MetaLearning, and is central to the mitigation strategies employed to alleviate the identified barriers, discussed in Section 7.2. This section addresses RQ3, RO2.b, RO2.c and RO2.d. This section includes a discussion on best practice relating to online pedagogy, a review of MetaLearning and how effective users who trialled the tool found this. A framework of topics for an introductory course is also discussed based upon the findings from this research and the tutorials within MetaLearning.

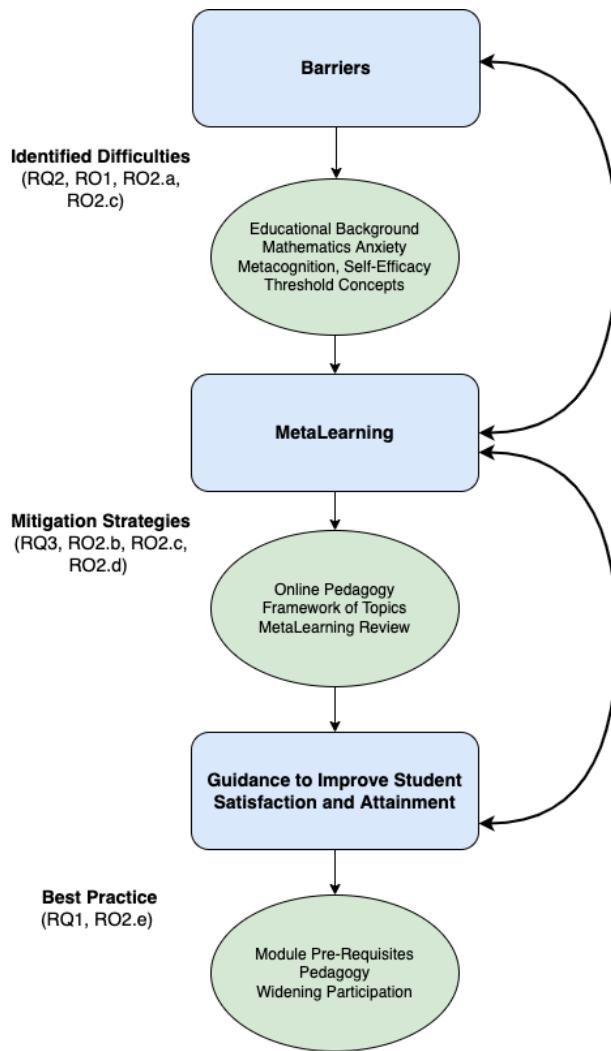


Figure 47: Diagrammatic Overview of Framework

Finally, Section 7.4 discusses all of the findings throughout the differing stages of this research and the previous elements of the framework to identify how this research can be used by others who teach some form of AI to improve student satisfaction and attainment (RO2.e). A summary is provided at the end of this chapter outlining how this framework can be considered as potential best practice for teaching this domain.

7.2 Barriers to Learning Machine Learning

As outlined in Chapter 1, there is currently a lack of research pertaining to the best practices for teaching Machine Learning, a domain which requires substantial knowledge and skills in both mathematics and computing. Therefore, this research endeavoured to identify the potential barriers and obstacles to learning (RQ2) as well as examine specific strategies to

assist learners in overcoming these barriers (RQ3). These aims were outlined in the following research objectives (*RO2.b, RO2.d and RO2.e are addressed further on in this chapter*):

1. Identify and examine the barriers that might impact upon student attainment in ML courses, using Machine Learning modules at Newcastle University and other participating institutions.
2. a. Identify the threshold concepts within Machine Learning
2. c. Improve student's metacognition and self-regulation regarding their learning of Machine Learning through implementation of strategies such as testing as a learning tool and knowledge surveys within the learning resource.

7.2.1 Educational Background and Mathematics Anxiety

As outlined in Chapter 2 (p.9) within the literature review, Machine Learning is primarily taught within Computer Science departments, where the focus is usually on teaching data analysis tools and the corresponding programming languages and, rarely the mathematics and statistics underlying this domain (Deisenroth, Faisal and Ong, 2020). This lack of core attention alongside concerning findings from the UK data skills gap investigation (Department for Digital Media and Sport, 2021), that the UK does not currently place quantitative skills at the core of the curriculum led to the conjecture that a lack of mathematics skills, potentially leading to mathematics anxiety may be a barrier to learning in Machine Learning courses.

The study, discussed in Chapter 4 (p.81) outlined the current experiences of AI within HE which encompassed both questionnaires and interviews with lecturers experienced in teaching this domain. The majority of respondents identified that their students encountered issues with mathematics, particularly the theoretical aspects of this domain (Sections 4.4.1, 4.4.2, p.96). This finding contributes to the resolution of RQ2 (p.6) in that mathematics and the theoretical aspects of this domain pose a problem to students. Alongside the difficulty with the theoretical aspects, the lecturers also noted that some students have difficulty grasping the terminology. As discussed in the literature review, the potential variability in student background has been linked to motivational constructs such as self-efficacy and mathematics anxiety (Lee, 2009). This variability in educational background was also identified within the lecturer interviews (Section 6.5.3, p.197) where all participants highlighted the differing level of prior mathematics understanding of their students. One of the participating lecturers also

advised that their students often drop out of the module when they realise the amount of mathematics involved so this may have implications for widening participation in the field.

The case studies discussed in Chapter 4 (Sections 4.5.1, 4.5.2, 4.5.3) investigated the level of prior mathematics attainment as well as the confidence level of students with their mathematics skills. Within the case study, the two universities that participated in the questionnaires had a differing average level of attainment. The two sets of students had differing educational background as one module was postgraduate and the other undergraduate. However, the students were specifically asked about their highest level of mathematics attainment. The majority of students on the postgraduate module had A-Level Mathematics whereas the majority of students on the undergraduate module had GCSE Mathematics. Although the students on the postgraduate module had the higher mathematics attainment, the students on the undergraduate module had a higher reported confidence level in their mathematics skills. This disparity between mathematics attainment and confidence may be a reflection of the students' metacognitive ability, as this can be affected by either overconfidence or mathematics anxiety. Therefore, lack of metacognitive skill may be a barrier to learning (RQ2). The cognitive mitigation strategies within MetaLearning may assist learners in overcoming this difficulty (RQ3).

Within the pre-module questionnaire, the postgraduate students expected to find the theoretical aspect of the module harder than the practical aspect, potentially indicating that the students initially felt greater confidence in their programming skills as opposed to their mathematics. However, this would depend on the learner's mental model of the domain, as they may not be cognisant of the foundational element mathematics and statistics plays within the AI domain. This lack of mental model and comprehension of the role of mathematics was highlighted by Lecturer B in the lecturer interviews (Section 6.5.3, p.197) and by the respondent to the student questionnaire relating to the online tool (in section 6.4.3, p.190) who responded that the Maths and Stats tutorial for Machine Learning was "*completely unnecessary.*" Therefore, an absence of an accurate mental model may be a barrier to learning (RQ2) due to a fundamental lack of understanding of what the domain entails.

All of the modules within the case studies contained specific mathematics and statistics content. Both the postgraduate Machine Learning module (Section 4.5.1, p.101) and the undergraduate Machine Learning and Computer Vision module (Section 4.5.3, p.118) offered a session which aimed to be a “primer” or “refresher” of mathematics and statistics specific to Machine Learning. These sessions covered content such as differential calculus, probability, vector and matrix algebra. These topics were covered in a single session. Observation of these sessions identified that these concepts were introduced at a high level meaning that students were expected to have some prior understanding of these concepts due to the fast pace of the delivery and, particular to the postgraduate module, comments from the lecturer that this material was “simple” and “straightforward” (Section 4.5.1, p.109). These comments may potentially negatively affect a student’s self-efficacy and confidence pertaining to mathematics and statistics if they found the content of these sessions complex. Across all of the modules within the case studies, the students regularly asked questions pertaining to the mathematical notation on the lecture slides, indicating a lack of familiarity with these terms. Although comprehending the mathematical foundations of this domain is important, as identified by Hicks and Irizarry (2017) minimisation of mathematical notation is a recommended best practice for teaching within this domain. Within Machine Learning instruction it may be beneficial to include a visualisation or textural explanation of a model/algorithm alongside the mathematical definition. The use of visualisation and a textural description may be more intuitive and assist the student in understanding the mathematical notation (RQ1).

The variation in educational background and identified issues with the mathematical elements of this domain by the lecturers, highlighted the importance of inclusion of a tutorial within MetaLearning which covered the mathematics and statistics knowledge pertinent to Machine Learning. In order to align and provide a baseline level for student mathematics knowledge, the tutorials explained the concepts at a level which required a low level of prior knowledge. However, the results from the student questionnaires reviewing the online tool (Section 6.4.3, p.179) found that the majority of students found the Maths and Stats for Machine Learning tutorial the most difficult out of all the tutorials included within MetaLearning. The majority of students noted that this tutorial had improved their knowledge in this area, although a proportion of the respondents (31.2%) indicated that they were not sure, potentially

indicating an issue with their metacognition as they could not accurately reflect on any changes to their skill level. Inclusion of more contextualised tutorial material so that the link between the mathematical content and Machine Learning is more explicitly stated may be beneficial to include within MetaLearning. As well as greater opportunity for users to reflect and assess their understanding to further boost their metacognition.

Particular aspects of the Maths and Stats tutorial that the students had the greatest difficulty understanding related to either probability or statistics. This finding was identified by the responses to the knowledge survey (Section 6.4.4, p.190). This may indicate a particular issue relating to the current educational provision and pedagogy relating to statistical instruction. Therefore, there may be a greater need for better quality instructional materials and teaching of probability and statistics.

The process of identifying the barriers to learning Machine Learning also encompassed the practical/programmatical aspects of this domain, to determine how the students perceived this element of learning the discipline. Students participating in the questionnaires as part of the case study were asked to rate their skills as a programmer and whether they identified as a “beginner”, “novice” or “expert.” The most frequent response for both cohorts of students was “novice”, however female students on the postgraduate Machine Learning module were more likely to describe themselves as “beginners” as opposed to their male counterparts. This could potentially indicate a lack of confidence or self-efficacy in their programming skills. The participants were also asked to rate their confidence in their programming skills on a scale from 1 to 10, with 10 being exceptionally confident. Students on the postgraduate and undergraduate modules had a very similar average confidence level with a confidence level of 6.2 and 6.1 out of 10 respectively. This similarity in confidence level relating to programming potentially indicates a lack of self-efficacy relating to the students on the postgraduate module as the majority had an undergraduate degree in Computing. Both student cohorts identified the practical aspects as harder in the post-module questionnaires. The students who reviewed MetaLearning rated their confidence in applying their knowledge in a practical context as a 3 out of 5 and there was variation in the responses with some students placing their confidence below a 3. Therefore, greater attention and instruction should also be placed on the practical/programming aspects of modules within this domain to ensure students are progressing in both the practical and theoretical aspects.

The variation in educational background and identified difficulties relating to both the theoretical and practical aspects of Machine Learning potentially indicate difficulties for students in combining these two fundamental aspects of knowledge pertaining to mastery of this domain. This issue was highlighted by one of the lecturers interviewed (Section 4.4.2, p.100) who discussed how their students struggled to match an appropriate algorithm to the dataset. This disconnect between the theoretical and practical, indicates the importance of an active learning pedagogical approach to this discipline so that the learner can undertake a number of tasks to build their understanding and mental model and construct these actions into further levels of functioning and understanding (RQ1). However, the student requires a sufficient level of surface understanding before they can link both the factual and procedural knowledge as outlined by Krathwohl (2002). Therefore, the inclusion of problem-based and active learning through the use of real-life scenarios and programming exercises may be being introduced too soon in these types of modules leading to the difficulty identified in connecting both the theoretical and practical knowledge. This connection and difficulty interleaving the theoretical and practical aspects of the domain may be a barrier to learning (RQ2).

7.2.2 Metacognition, Self-Efficacy and Self-Regulation

As discussed in the previous section there were a number of findings pertaining to educational background and mathematics anxiety which indicated potential issues with student self-efficacy and metacognition including the relatively low confidence levels pertaining to both mathematical and programming skills. As identified in the literature review (Section 2.8.2, p.51), students who have low self-efficacy are more likely to suffer from low confidence in their skills. Intrinsic to the concept of self-efficacy is metacognition, particularly the learner's ability to accurately reflect and self-regulate to understand their own learning. Equipping learners with the skills and tools to assist them in becoming more self-regulated was the aim of RO2.c (p.6).

As identified in the literature review, there are a variety of mitigation strategies which can be employed to improve self-efficacy such as self-directed mastery opportunities recommended by Bandura (1977) and independent performance. The literature guided both the creation of the online tool and the implementation of the mitigation strategies.

Students were surveyed on a number of elements pertaining to their self-efficacy and metacognition, within the pre-module survey. Students on the postgraduate Machine Learning module were asked on a scale of 1-10 (with 10 being exceptionally confident) how confident they were in their ability to do well in the module. The results for this question were mixed as some students answered as low as a 2 or 3, the majority of students rated their confidence as a 5 or 6, this does not indicate that the students have a strong sense of self-efficacy in their ability to do well. This questionnaire was completed by students after the first lecture of the module, a potential extension to this research would be to survey students on their confidence level before undertaking the first lecture and then after, as this may indicate an incomplete mental model of the domain which may be affecting student confidence in their ability to pass the module. This may also relate to the comment made by lecturer B (Section 6.5.3, p.197) who advised that some students drop out when they realise the amount of mathematics involved. Therefore, incorporating strategies to improve student self-efficacy in this area may ease an element of anxiety and lack of confidence (RQ3).

RO2.b, the creation of the learning resource, enabled the inclusion of mitigation strategies and in itself became a strategy as distance learning has been proven to lessen student anxiety relating to mathematics (Taylor and Mohr, 2001) (RQ3). One of the participants in the UKICER workshop (Section 6.6.2, p.204) mentioned that the classroom environment does not work for all students. Therefore, offering MetaLearning as an alternative mode of instruction may be beneficial to these particular students so that they have differing learning opportunities available to them.

Within MetaLearning, two specific methods were employed to help improve self-efficacy and metacognition, these were identified in RO2.c (p.6) as knowledge surveys and testing as a learning tool. The aim of the knowledge survey was to enable the learners to reflect on their current knowledge and confidence level pertaining to their mathematics and statistics knowledge, completing this would signpost to students' current gaps in their understanding and aspects which require further attention. Upon completion of the Maths and Stats tutorial, the students were encouraged to complete the knowledge survey again to help them identify areas where they had improved both their understanding and their confidence to boost their self-efficacy and reduce their mathematics anxiety. From the students who trialled the tool, the majority responded to the questionnaire that they felt the knowledge survey has helped

them reflect on their mathematics and statistics knowledge. Therefore, this opportunity to reflect may assist the students in building more comprehensive metacognitive strategies, allowing them to more accurately reflect on their skill level, boosting confidence and self-efficacy (RQ3). The lecturers who reviewed the tool, through the questionnaires and workshop, also felt that the knowledge survey would be beneficial to improve confidence. However, one respondent felt that this technique would be best suited to a blended learning environment where the learner completed the knowledge survey on their own, then was supported in their learning of the areas they were unsure on by a lecturer or tutor. MetaLearning was created to enable deployment as both a standalone resource and for use within blended learning, therefore use of the knowledge survey within either context is supported.

Testing as a learning tool consisted of low stakes questioning and the provision of feedback to enable the users to have mastery experiences which have shown to improve self-efficacy. Within the UKICER (2021) workshop, the importance of quality feedback was highlighted by participants as a key method for improving confidence and self-efficacy (p.204). Although MetaLearning provides users with feedback on whether they answered the question correctly and provides guidance on the pass mark, further research and an extension of the feedback provided through the tool would be beneficial to further the current provision. All of the participants who reviewed the online tool were positive about the benefits of low stakes quizzing, the majority of students found this mitigation strategy useful (RQ3). The students who trialled the tool also felt that overall, it had led to an improvement in their confidence in their AI and Machine Learning knowledge (Section 6.4.3, p.179). Therefore, testing as a learning tool has proven a useful learning strategy to not only assist learners in comprehending the material but also as a method to improve confidence in their knowledge (RQ3).

The variability in students' educational background discussed in Section 7.2.2, both societal and educational, may impact upon their ability to employ appropriate learning strategies to overcome these difficulties, for example by linking their previous learning with new material. There were clear differences in the learning strategies employed by students within the case studies. The students on the postgraduate Machine Learning module indicated that they would not use more holistic learning strategies such as perseverance, resilience, goal setting

and reflection. However, the students on the undergraduate Artificial Intelligence module indicated that they used both goal setting and reflection as learning strategies. Both cohorts of students were asked to rate their confidence from 1 to 5 relating to how confident they were in applying the ML/AI techniques upon module completion. The responses from the postgraduate module were varied with participants rating themselves on every point of the scale apart from a 1 (low confidence). Within the undergraduate module, responses were less varied with the majority of students placing themselves as a 3 on the scale. Determining the causation of the differentiation in confidence on the case study modules has potential as a continuation of this research to determine the efficacy of certain learning strategies on student confidence.

If students have low self-efficacy, they may not have the emotional capability to persist when in the liminal state of a threshold concept. Equipping students with metacognitive skills and training them to become more self-regulated learners may help build up student resilience when encountering challenging educational scenarios.

7.2.3 Threshold Concepts and Pedagogical Content Knowledge

Within the literature review (Section 2.7, p.44) the concept of pedagogical content knowledge (PCK) was outlined by Shulman (2013) as the idea that equal attention should be paid to the content components of teaching as is usually devoted to the teaching strategy. This includes the preconceptions of the student and how to effectively describe ideas and difficult topics to teach including the threshold concepts within this domain. As identified in Chapter 4 (Section 4.3.6, p.94) the variation in modules offered within this domain and the disparity in content covered under similarly titled modules, for example Machine Learning modules, suggests that the PCK has not yet been clearly outlined.

Chapter 4 outlined the current provision and experiences within HE, from the online review of modules (Section 4.3, p.82). The most common topics taught across Machine Learning, Deep Learning and Artificial Intelligence modules included specific types of algorithms including classification, regression and neural networks and specific algorithms consisting of Naïve Bayes classifier (Gandhi, 2018), Support Vector Machine (SVM) (Cortes and Vapnik, 1995), Decision Trees (Quinlan, 1986), Convolutional Neural Networks (CNN) (LeCun *et al.*,

1998) and Recurrent Neural Networks (RNN) (Hochreiter and Urgen Schmidhuber, 1997). The sub-domain of Deep Learning was prominent across all modules.

Several data collection methods were used to start to identify the threshold concepts within this domain, including post-module questionnaires for institutions participating in the case study and the one-minute paper. Students studying on the postgraduate Machine Learning module completed two rounds of the one-minute paper (Section 4.5.1, p.111) to determine which topics from this module they identified as troublesome, these included SVM, Multi-Layer Perceptron (MLP) (Brownlee, 2016), RNN, CNN and specific domain applications including Deep Learning for human activity recognition and Machine Learning for computer vision. The identification of the domain applications as troublesome potentially indicates a disconnect between the specific techniques and application to a practical domain. The one-minute paper participants also identified backpropagation (Rumelhart, Hinton and Williams, 1986) and feature engineering as areas of the module which they didn't fully understand. This same cohort of students completed a post module questionnaire in which they were asked to identify which were the most difficult topics in this module, the most frequent responses were concepts relating to Deep Learning as well as backpropagation. There was some disparity and split between students who did and did not find the mathematical aspects difficult. The post module questionnaire completed by participants within the undergraduate AI module identified the K-means algorithm (Bock, 2008), neural networks, MLP and backpropagation as the most difficult concepts they learnt within the module. The commonality in which these topics arose not only assisted in identifying potential threshold concepts but also as difficulties students face when learning this topic (RQ2).

The students who reviewed MetaLearning were asked which of the tutorials they found the most difficult. The most frequent responses included the Maths and Stats tutorial, Machine Learning Algorithms and the Overview of Deep Learning tutorial. All of the tutorials correlate to the findings from the case studies in that students have difficulty with certain Machine Learning algorithms and have particular trouble with the sub-domain of Deep Learning. However, the finding that most students found the Maths and Stats tutorial the most difficult is a strong indicator that there is a potential issue understanding the theoretical underpinnings of this domain.

Alongside the data collected from students, questionnaires (Section 4.4.1, p.96) and interviews (Section 4.4.2, p.97) were conducted with lecturers to determine which areas of this domain they noted that students struggled with. Specific points raised by the lecturers included the observation that students can become overwhelmed in lectures, however students with a mathematics and statistics background often find it easier. Another participant identified that students often had difficulty identifying the terminology and matching an appropriate algorithm to a dataset. These findings provided motivation for MetaLearning to try and ensure that students have the mathematical basis which may stop them being overwhelmed in lectures, as well as helping them get to grips with the terminology. Providing the students with a baseline, introductory understanding should help the students gain a fundamental understanding of the theory of this domain to help them apply this knowledge within a practical context (RQ1).

The findings from Section 4.6.4 (p.126) signified that AI and Machine Learning is an inherently active learning discipline, comparable to other Computer Science disciplines. Therefore, learning strategies inherent to the PCK include group activities, questioning and problem-solving opportunities. Within the case studies these led to increased engagement with lecturers and led to students being more forthcoming with questions and interaction with their peers (RQ1).

7.3 Online Resource for Introduction to Machine Learning (MetaLearning)

Within the research objectives outlined in Chapter 1 (p.6), there were two objectives pertaining to creation of the learning resource and a collection of topics deemed fundamental to an introductory Machine Learning course (RO2.b, RO2.d). Creation of MetaLearning also related to all three research questions as mitigation strategies were included within the resource (RQ3), and the review by both students and lecturers would provide further data pertaining to difficulties faced when learning this domain (RQ2) and potential best educational practice (RQ1).

Findings within the literature review (Chapter 2) indicated that hosting the learning resource online may be beneficial to improve learner self-monitoring, and therefore their metacognition especially through the use of quizzing (Section 2.9, p.55) which was identified as a strategy in RO2.c. As discussed in Section 2.6.2 (p.44) a study by Schwab-McCoy, Baker

and Gasper (2021) which surveyed lecturers teaching Data Science found that respondents felt there was a need for more teaching resources, specifically online tools which would enable them to further embed active learning within their teaching. Alongside the findings from Chapter 2, the lecturer questionnaires and interviews in Chapter 4 (Section 4.4.1, 4.4.2) identified that online resources were the most popular resource for additional learning materials. Students responding to the post-module questionnaires (Sections 4.5.1, 4.5.2) also identified that they used online resources for additional support. These findings guided the decision to create an online learning tool.

7.3.1 *Online Pedagogy*

Determining effective pedagogy for the online tool was key to ensure the system was designed as effectively as possible. As outlined in Chapter 2 (Section 2.9.2, p.57) a constructivist approach to instructional design for online learning is advantageous to learning within this situational environment (Oliver, 2001). Specific strategies relevant to the constructivist approach including contextualising knowledge, motivating self-awareness and multiple modes of representation were included within the online tool (RQ1). The use of visualisation and contextualisation is also recommended best practice for teaching within the Data Science domain (Hicks and Irizarry, 2017), therefore inclusion of these strategies should assist students in learning the MetaLearning material (RQ1).

MetaLearning was designed with consideration of the Hattie and Donoghue (2018) learning model (Section 2.5.3, p.32), for example surface learning was facilitated through inclusion of specific subject matter vocabulary and encouragement to take notes. The consolidation phase of learning is encouraged through practice testing through the use of low stakes quizzing and feedback (Hattie and Donoghue, 2016). The progression through the different learning stages is part of the educational deep learning process. Strategies to encourage this process included within MetaLearning consist of self-monitoring, reflection and evaluation. These techniques are encouraged and promoted through use of the low-stakes quizzes which enable the learners to monitor specific topics which they are competent in as well as the use of knowledge surveys which assist the learners in reflecting on their mathematics and statistics knowledge.

Developing a learning resource which endeavours to alleviate a learners mathematics anxiety and to improve their self-efficacy was a key motivator for the online tool as Section 2.8.1 (p.47) discussed, retesting and self-paced learning have been shown to reduce mathematics anxiety. The mitigation strategies incorporated within the tool aimed to alleviate some of the potential barriers discussed in the previous Sections (7.2.2 and 7.2.3) with the overall aim of addressing RQ3. The knowledge survey was deployed as a tool for reflection to help learners build a stronger metacognitive understanding of their mathematics and statistics knowledge and to potentially build their confidence and self-efficacy upon completion of the tutorial. The use of testing as a learning tool, specifically low stakes quizzing aimed to provide the learners with self-directed mastery opportunities which have been shown to enhance self-efficacy (Section 2.7.2). Hosting the learning resource online is in itself a mitigation strategy as it allows for self-paced distance learning.

7.3.2 Framework of Topics for an Introductory Machine Learning Course

Outlined within the research objectives in Chapter 1 (p.6), one of the aims of this research was to create a framework of topics for an introductory Machine Learning course (RO2.d). The initial motivation for this objective was the lack of current research pertaining to education within this domain and the rising demand for graduates skilled within AI and specifically Machine Learning. RO2.d also helps address the research question pertaining to best practice (RQ1) as the framework of topics will provide practitioners with an overview of what is considered best practice to teach within such a module. As discussed in the literature review (Section 2.6.2, p.40) the report on the UK data skills gap (Department for Digital Media and Sport, 2021) identified specific skills which potential employers specify that graduates are lacking, these include Machine Learning, data processing and data ethics. Therefore, it was pertinent to design the content of MetaLearning around these skills.

One of the key findings from the literature review was that lecturers within this domain have great difficulty in narrowing down the vast amount of content into an introductory course (Schwab-McCoy, Baker and Gasper, 2021). This finding from the literature was echoed in the findings from Chapter 4 by both lecturers and students. Respondents from the lecturer questionnaires (Section 4.4.1, p.96) indicated that students often claim the pace of the module is too fast and that students often complain about the wealth of material. The lecturer of the AI module within the case study (Section 4.5.2, p.115) also corroborated this view in

that there is a wealth of potential content to be fit within a module. The development of the PCK for the domain should assist the lecturers when creating their modules to narrow the scope of content to be included, particularly within an introductory course. As well as causing difficulty for the lecturers, the wealth of content was also raised by students as a potential barrier to learning (Section 4.5.1, p.110). Within the postgraduate Machine Learning module (Section 4.5.1, p.110) one student advised that the amount of content to learn was much greater compared to other modules they were undertaking, another felt that it was difficult to 'digest' all the information and that they were having difficulty keeping up with the material. Therefore, working towards identification of key topics should narrow down the scope and disparity in offerings within the educational provision of this topic.

The findings from Chapter 4 were key to the formulation of a framework of topics and these findings were central to the implementation of the online tool in determining the content for inclusion in the tutorials. The online review of modules (Section 4.3, p.82) identified topics which were most frequently taught on Machine Learning, Deep Learning, AI and Data Science modules to determine if there were a core set of topics germane to these disciplines. The findings from the review of the Machine Learning and Deep Learning modules were the most pertinent towards the creation of the framework of topics for an introductory Machine Learning course. Within the review of the Machine Learning modules (Section 4.3.2, p.84) the most commonly taught topics included classification, regression, clustering and neural networks. Specific algorithms included Naïve Bayes classifier (Gandhi, 2018), Support Vector Machine (Cortes and Vapnik, 1995) and Decision Trees (Quinlan, 1986). Deep Learning was also taught on the majority of Machine Learning modules including content on Convolutional Neural Networks (LeCun *et al.*, 1998) and backpropagation (Rumelhart, Hinton and Williams, 1986). A concerning finding from the review of Machine Learning modules was that only 10% taught some form of ethics or content relating to the legal and social issues surrounding the use of Machine Learning discussed in Chapter 2 (Section 2.3, p.19). The review of the Deep Learning modules identified that the most commonly taught topics included Convolutional Neural Networks (LeCun *et al.*, 1998), Recurrent Neural Networks (Hochreiter and Urgen Schmidhuber, 1997) and Generative Adversarial Networks (Goodfellow *et al.*, 2020).

The questionnaires and interviews with lecturers detailed in Chapter 4 (Sections 4.4.1, 4.4.2) also identified the most commonly taught topics, listed in Table 26, including linear and logistic

regression (Worster, Fan and Ismaila, 2007), K-Means (Bock, 2008) and Deep Learning algorithms such as CNNs and RNNs. Within the lecturer interviews, the participants were asked which topics they feel are pivotal topics to teach, respondents identified Deep Learning, in particular CNNs and LSTMs (Brownlee, 2017).

Machine Learning	Deep Learning
Linear and logistic regression	Convolutional Neural Networks (CNN)
K-Means	Recurrent Neural Networks (RNN)
Principal Component Analysis (PCA)	Backpropagation
Support Vector Machine (SVM)	Long Short-Term Memory (LSTM)
Bayesian Machine Learning	

Table 26: Topics mentioned in lecturer questionnaires and interviews

The institutions participating in the case studies (Section 4.5.4, p.119) all had some form of similarity in content. They all taught supervised and unsupervised learning as well as Deep Learning. Specific domain applications were also covered on all three modules, particularly computer vision. Inclusion of a specific field within AI enables the students to comprehend the breadth of application scenarios of the differing algorithms they are learning as well as the scope of achievable outcomes.

Table 27 outlines the key topics identified as most frequently taught within this research, the table is categorised into differing types of algorithm including supervised and unsupervised and artificial neural networks. Depending on the type of module and the situational organisation, the content in Table 27 may still be a large amount to teach within a condensed timeframe.

Algorithm Types	Algorithms
Supervised Learning – <i>Classification, Regression</i>	Naïve Bayes, SVM, Decision Trees, Linear Regression, Logistic Regression
Unsupervised Learning – <i>Clustering, Dimensionality Reduction</i>	PCA, K-means
Artificial Neural Networks	CNN, RNN, LSTM, GANs

Table 27: Topics identified as most commonly taught within this domain

It should also be noted when considering Table 27 as a potential framework for an introductory course for Machine Learning that a number of the algorithms were also identified as potential threshold concepts, including SVM, Decision Trees, CNNs and RNNs. Therefore,

potentially more time will need to be dedicated to the teaching of these topics. The majority of content included within Table 27 is covered within the tutorials in MetaLearning as these topics were identified as key to teach within this domain.

7.3.3 *MetaLearning Review*

Determining how useful both lecturers and students perceived MetaLearning to be, as well as the integrated mitigation strategies was imperative to comprehend how effective the resource is as a learning tool. Both lecturers and students gave the tool a positive review, 92.2% of the students said that the tool had improved their understanding of Machine Learning (Section 6.4.3, p.181) and the majority of participants said they would reuse the tool. All the lecturers interviewed felt that the tool would help their student's understanding of Machine Learning (Section 6.5.3, p.197). The majority of participants at the UKICER (2021) workshop also felt that MetaLearning would help with student understanding of this domain. The majority of the lecturers rated a '5' on the scale of how likely they were to recommend the tool to their users, they also felt that the content was pitched at the right educational level for their students.

The students who reviewed MetaLearning found the visualisation of specific concepts particularly useful, 89.61% of respondents felt that it has aided their learning (RQ1). This is potentially a technique to extend within the classroom to assist lecturers in explaining difficult concepts, it could also be deployed to assist with the learning of the mathematics and statistical elements of this domain. This technique could also be applied wider within the online tutorials, especially within the Maths and Stats tutorial, however reviews of this tutorial were positive from both lecturers and students.

7.4 Guidance to Aim to Improve Student Satisfaction and Attainment

One of the main aims of this research was to investigate potential best practice and as outlined in RO2.e (p.6) to discover ways of improving student satisfaction and attainment within these courses. Retaining students within these modules may also be a potential issue as outlined by Lecturer B (Section 6.5.3, p.197) if the students do not have an accurate mental model of what this domain entails. Determining the current demographics of students studying within the domain will also instruct on further research relating to widening participation.

7.4.1 Module Pre-Requisites and Pedagogy

Determining at what educational level these types of modules are offered can indicate potentially the level of complexity of the material taught on this module and for example the types of learning strategies and coping mechanisms the students may employ, as postgraduate taught students will have more experience learning within HE than undergraduate students. However, findings from the case studies and particularly from the modules who participated in the questionnaires, indicated that the undergraduate students utilised a greater variety of study strategies, including reflection and goal setting compared to the postgraduate students. Therefore, incorporation of strategies such as the knowledge survey and low stakes quizzing will help promote these types of skills for students who don't use these learning techniques.

From the online review of modules (Section 4.3, p.82) the majority of the modules offered, which encompassed some form of AI were offered at undergraduate level, apart from Deep Learning which was offered at both undergraduate and postgraduate. Pre-requisites for these modules were similar and were mainly centred around mathematics and statistics knowledge and some form of programming experience. However, it was not clear from the module descriptors how conformity to these pre-requisites were assessed. There are both positives and negatives to the outlining of pre-requisites, it enables the module leaders to outline a baseline understanding of the foundational topics they require students to have to help them understand the course content. However, they can be a barrier to widening participation, especially if the students are not particularly confident in their mathematics and programming skills. Therefore, due consideration should be made before outlining module pre-requisites to determine if they are strictly necessary and whether students may struggle with the module content if they do not have a thorough grasp of foundational knowledge in which the content will build upon.

The lack of prior mathematics knowledge identified by the lecturers may potentially exacerbate any issues the students may have with the programming aspects of the module which both cohorts of students within the case study found the most difficult (Section 4.5.4, p.120). For example, if students do not understand how to decide on an appropriate algorithm given a dataset, then this may affect their confidence in their practical abilities. Two of the

lecturers interviewed when reviewing MetaLearning (Section 6.5.3, p.199) mentioned a potential need for a specific programming tutorial within the online tool. This will be considered for further extension and iteration of MetaLearning. However, as indicated within the literature review (Section 2.6.2, p.41) it may be of benefit for the students to explore the idea of Machine Learning in general before delving into the programming syntax as this may help the learners build a clearer mental model of this domain (RQ1).

All of the modules participating in the case studies and the majority of modules examined as part of the online review of modules were taught through a mixture of lectures and practical sessions. This is comparable with other modules taught within Computer Science programmes in that the main theoretical ideas are taught within the lectures, and within the practical sessions students are supported through practical tasks which help embed the theory and support the students to develop their practical programming skills.

There were a number of pedagogical techniques identified through the case studies, interviews and questionnaires with lecturers and the review of MetaLearning which may constitute best practice for teaching this domain (RQ1). Within any educational context, awareness of student engagement levels and the variation of skills and understanding a student possesses is incredibly important. It is a valuable technique that lecturers are adaptable within their learning methods they employ to meet the varying student requirements. Aligning with the threshold concepts and recognising that some topics will require a greater amount of time to be spent ensuring students understand these concepts is particularly relevant within the field of AI as there are a range of foundational issues such as feature engineering and backpropagation which are often difficult to comprehend. These topics are essential to comprehend in order to build a Machine Learning model and were identified within this research as topics students had difficulty with. One of the most prominent strategies to assist students in learning this domain and helping them to overcome the threshold concepts was the use of real-world examples. These were used to facilitate and to contextualise the theoretical information the students were learning. For example, Deep Learning, especially models categorised as artificial neural networks (ANNs) (Abiodun *et al.*, 2018) were identified as a potential threshold concept, one method employed within the case study AI module (Section 4.5.2, p.116) to assist students in building a more comprehensive mental model was to compare ANNs to human learning and neural networks in the brain.

Students are often engaged relating to learning about themselves and this correlation may assist them to build a greater mental model of Deep Learning. However, it is essential to indicate that this isn't an exact mapping between ANNs and human learning to avoid additional confusion (RQ1). Students who lack a clear mental model of the domain may have low self-efficacy and confidence in their belief and competence in the module. Students without an effective mental model may be "susceptible to the fiction that ML has a "hidden mind" (Fiebrink, 2019). Therefore, it is of particular importance within AI modules to clearly define key concepts.

Due to the wealth of content which is often included within these modules as previously discussed in Section 7.3.2, it may be beneficial when introducing new models or algorithms to link these back to previous content covered and discuss in anyway which they are similar or dissimilar. This was a strategy employed on the postgraduate Machine Learning module (Section 4.5.1, p.110) where a recap was provided at the beginning of each lecture as well as a comparator to provide baseline information on any commonalities or understanding/knowledge from previous models which are pertinent to the new topic. Within the Machine Learning module (Section 4.5.1, p.110) students were advised to review their theoretical understanding of a specific algorithm after they had implemented/coded it within the practical session to solidify their theoretical understanding. However, this strategy may assist students who struggle with either aspect of the module as this iteration of revision of the content should help facilitate learning in both the theoretical and practical elements.

As discussed in Section 7.2.2, minimisation of use of mathematical notation in lecture slides should be limited. However, this does not mean that it should be eliminated completely as it is important that students have an understanding of the foundations of Machine Learning. There was an element of disparity in how the lecturers in the case studies handled the mathematical elements of their modules, within the postgraduate Machine Learning module (Section 4.5.1, p.110) presentation slides were heavy with mathematical notation and were used as the main source of instruction to explain the algorithms, causing students who did not have a particularly strong mathematics background difficulty. However, the lecturer for the undergraduate Machine Learning and Computer Vision module (Section 4.5.3, p.119) used mathematical notation sparingly, dedicating time to explain what the notation meant/represented and asked the students questions on it to ensure they understood.

Therefore, it may be beneficial to include this content within lectures so that students comprehend the foundational elements of the domain, but it should be used when strictly necessary and not as the main instructional method to explain a new algorithm (RQ1).

A beneficial technique to scaffold engagement in lectures, recognised within the observations of the institutions participating in the case studies, was the use of group activities. Group activities have been shown to increase motivation as well as offering the students a perceived support system which is often not felt when working individually (McLean, 2009; Lavy, 2017). Initiating the lectures by giving an overview of the session structure and informing the students that there will be an activity led to students paying greater attention and displaying more engagement with the content. It also appeared to foster a more dynamic and open relationship not only between peers, but also between the students and the lecturer. As a consequence, there was a notable difference between the number of questions the students asked the lecturer in the module with the group activities as opposed to the one that did not offer these (RQ1).

Assessment for modules within the field of AI were similar for each of the sub-fields including Machine Learning and Deep Learning. The most popular form of assessment was a mixture of coursework and exam. This is an often-used assessment procedure as the mixture of assessment type can add variety to a learners experience, this variety may also ensure that all students can demonstrate their strengths within an assessment context they feel most comfortable with (Race, 2020). Within the AI domain, the theory is just as important as the practical skills, therefore the combination of the two assessment types will enable students to receive feedback pertaining to both of these knowledge/skill sets to identify areas which require further learning.

7.4.2 Widening Participation

The current impetus to widen participation within the AI domain is driven by a number of factors outlined within chapter 2 (Sections 2.2.6, 2.3 and 2.4) including the need to establish development teams who are more representative of society and to ensure that the potential employment opportunities currently offered within this sector and benefits associated with these are available to a diversity of individuals. As outlined in Section 2.2.6 (p.19), the HE

sector will play a key role with the transition to the fourth industrial revolution and facilitating high levels of AI literacy within society.

To determine a sample of the current demographics of individuals studying within the AI domain, information pertaining to age and gender was collected from the institutions participating in the case studies. The majority of students on these modules were male and under the age of thirty. Although these results are limited to the specific cohorts participating in this research project, the results are indicative of this domain, where on average 78% of professionals are male (World Economic Forum, 2018).

Students who trialled MetaLearning were beginners to this field and often had very little prior experience or interaction with this domain. However, they indicated on the questionnaire (Section 6.4.3, p.189) that they were motivated to continue studying AI and Machine Learning. Therefore, potentially more emphasis should be placed on ensuring accessibility and broadening participation in such courses as there appears to be interest and motivation to learn this topic. A potential method to widen participation includes greater use of blended learning (Yang and Cheng, 2018), as this flexibility allows learners to fit their studies around employment or caring responsibilities. A tool, like the one created within this research could be utilised within a blended learning environment to build student's foundational knowledge of the domain. MetaLearning could also be used to upskill and improve AI literacy for individuals who were wanting an introduction to the AI domain.

7.5 Summary

In this chapter an overview of the potential barriers students may face when undertaking a course within this domain were discussed including mathematics anxiety, low self-efficacy and the threshold concepts within this domain (RQ1). To mitigate against these barriers, an online tool for Machine Learning was created, MetaLearning, and this chapter discussed the pedagogy, framework of topics for an introductory course and the lecturer and student review of the tool (RQ3). Finally, overall guidance relating to best practice for education within this domain was discussed with the aim of improving student satisfaction and attainment in AI modules (RQ2). The following chapter concludes this thesis in which the outcomes of this research will be summarised, and discussion of further work will be included.

Chapter 8. Conclusion

8.1 Review of Research Aim and Questions

The overall aim of this research as outlined in Chapter 1 (p.5) was to determine the barriers students face when learning AI as well as the difficulties encountered by lecturers teaching this domain to initiate a framework of best practice for AI education. To ensure the research aim is achieved three research questions were outlined (Chapter 1, p.5)

RQ1: What is good practice relating to the teaching of AI?

RQ2: What are the current perceived difficulties experienced by both students and lecturers relating to AI?

RQ3: How do cognitive mitigation strategies alleviate any identified issues encountered by students learning this domain.

To answer these questions Chapter 2 of this thesis explored the pertinent literature to identify existing research relating to best practice for teaching AI (RQ1), the review also assisted in identifying any barriers to learning which have already been identified from previous studies such as the prevalence of mathematics anxiety or difficult concepts (RQ2). Determining effective cognitive strategies which have previously been employed within Computer Science education were also discussed within Chapter 2 to determine strategies for inclusion within this research (RQ3).

The study methodology as discussed in Chapter 3 was guided by the overall aim of the research, the research questions and objectives. Analysis of both quantitative and qualitative data and methodological triangulation ensured a variety of views and data sources to answer the research questions.

Chapter 4 provided insight into the current AI education provision, enabling identification of current best practice for teaching this domain (RQ1) through a systematic online review of modules and the varying data collection methods employed within the case studies. These forms of data collection also enabled identification of the difficulties faced by both students and lecturers within this domain (RQ2).

Cognitive mitigation strategies were incorporated into MetaLearning (Chapter 5) to enable situational inclusion of these strategies within subject specific learning. Chapter 6 details the

review of MetaLearning, including the perceived success of these strategies as evaluated by students, lecturers and Computing professionals.

All research questions are addressed in Chapter 7, with the findings discussed within the framework. The variety and depth of data collected from a range of sources ensured a degree of certainty when outlining the study findings. To ensure that the research questions were fully addressed a number of research objectives were defined enabling specific facets of the questions to be focussed on. The following section discusses how these objectives have been met.

8.2 Review of Research Objectives

Two main research objectives were outlined in Chapter 1 (p.6) of this thesis. RO1 pertained to the identification of the barriers which might impact upon attainment within Machine Learning modules. This objective was addressed through the literature review in Chapter 2 (p.9) which identified that mathematics anxiety may be a potential issue as well as low self-efficacy and confidence. To determine whether these perceived barriers were true and to determine any other potential issues, qualitative research was undertaken as detailed in Chapter 4 (p.81) in the form of case studies consisting of observation, interviews and questionnaires with both staff and students. Although RO1 set out to identify the barriers, work completed relating to the other objectives also helped further comprehend difficulties students were encountering when completing a course within this domain, including the review of the learning resource created for RO2.b.

RO2 consisted of five sub-objectives which pertained to the aim of discovering how to alleviate some of the perceived barriers identified through RO1. RO2.a. related to the identification of the threshold concepts within Machine Learning. Although there is currently a lack of research pertaining to best practice for teaching this domain there are some studies pertaining to Data Science which were reviewed in Chapter 2 (p.40). The online review of modules and case studies within Chapter 4 (p.81) enabled identification of the most commonly taught topics and understanding of topics which lecturers feel students struggle with as well as the student perspective of the complex topics. The review of MetaLearning detailed in Chapter 6 (p.172) determined which of the tutorials the students found the most difficult, the findings of which correlated with the outline of potential threshold concepts from Chapter 4 (p.128).

The creation of the online learning tool for Machine Learning (MetaLearning) was in fulfilment of RO2.b. The learning resource was identified as most effective if situated online as identified in Chapter 2 (p.50) as this can ease mathematics anxiety and offers learners the chance to be more self-directed in their learning therefore, relevant to RO1 and RO2.c. Included within MetaLearning were the mitigation strategies identified in RO2.c to improve students' metacognition and self-regulation. Testing as a learning tool was implemented through the use of low stakes quizzes included within the various tutorials. Knowledge surveys were incorporated within the Maths and Stats tutorial as this subject area was noted as a particular barrier to learning within this domain.

To ensure MetaLearning contained all of the relevant topics pertinent to an introductory course for this domain a framework of topics was created. The framework of topics attained RO2.d. The framework was based on findings from the literature review which identified skills which graduates are currently lacking within key areas of the domain as well as the findings from Chapter 4 (p.81) including the most commonly taught topics from the review of online modules and the case studies. The findings from the case studies and qualitative research with both the lecturers and students also identified topics which learners struggle with the most. This informed how the tutorials were designed, with visualisation and real-world examples included for the more difficult topics to learn.

The final research objective, RO2.e, was based on improving student attainment and satisfaction within these courses. To determine a set of best practices for teaching AI all of the findings from the literature review (Chapter 2, p.9) the review of AI educational provision (Chapter 4, p.81) and the review of the online tool (Chapter 6, p.172) were analysed to determine any commonalities which might constitute best practice, and which were demonstrated within the findings of this research to assist learners when studying this domain. This guidance was provided and discussed in Chapter 7 (section 7.4, p.226).

8.3 Research Contributions

This thesis contributes to a limited pool of research relating to the educational provision of AI and specifically Machine Learning. The research contributes to a number of areas pertaining to Machine Learning education including the barriers students encounter when studying this domain. The findings included a disparity in educational background, particularly relating to

mathematics and statistics knowledge which potentially contributed to the initial expectation that the theoretical aspects of the subject would be the most difficult to learn, as outlined in the case studies (Section 4.5, p.101). Although the students participating in the research initially expected the theory to be harder, within the post-module questionnaire both cohorts actually stated they found the practical aspects most difficult (Section 4.5.4, p.120). Thus, indicating a potential lack of an appropriate mental model of the domain and difficulty combining the two aspects of the domain. Potential issues relating to metacognition and self-efficacy were also identified within the literature review (Chapter 2, p.51) and indicated as an issue within the cohorts participating in the case studies. The students had a relatively low self-identified confidence level in both their mathematics and programming knowledge. The online tool for Machine Learning (MetaLearning) which was created to mitigate against the identified barriers includes strategies such as knowledge surveys and testing as a learning tool which were both deemed useful methods by both the lecturers and students who reviewed the tool (Chapter 7, p.226).

MetaLearning became in itself a mitigation strategy to improve self-efficacy and ease mathematics anxiety due to the associated benefits outlined in Chapter 2 and Chapter 4. To ensure that the content of the tutorials were appropriate for the educational level and covered the key topics for an introductory course within this domain, a framework of core topics were collated from the research undertaken in Chapter 4 (p.81). This framework of topics discussed in Section 7.3.2 (p.223) can be used to assist module leaders in determining the core content for their courses as this was identified as an issue when planning courses within this domain. Alongside the framework of topics, an initial outline of potential threshold concepts were also defined, identifying particular difficulty with the mathematics and statistics underpinning this domain and that the sub-domain of Deep Learning is a particular issue for students studying this area.

The culmination of the analysis from the varying data collection methodologies employed during this research was overall guidance on best practice for teaching within this domain, including potential practice for improving student satisfaction and attainment. Use of active learning strategies including group activities, questioning and problem-solving opportunities displayed increased engagement within the lectures and students being more forthcoming with questions for the lecturer. This increased communication between peers and the lecturer

may in turn increase student self-efficacy and their confidence in their understanding of the material.

Identified difficulties pertaining to both theoretical and practical aspects of the modules may indicate issues with the current practice relating to pre-requisites. As discussed in Chapter 4 (p.100), the majority of modules offered within the AI domain have some form of pre-requisite relating to mathematical ability or previous programming experience. However, the findings that students still have difficulty with these elements indicates that additional revision within the modules may be necessary. For example by providing students with additional learning resources such MetaLearning so that they can build their baseline understanding and mental model before delving into the intricacies of this domain such as specific programming syntax. It may also be beneficial to further investigate the use of module pre-requisites to determine if they are effective and to understand any impact they can have on widening participation, a particular issue within the AI domain.

8.4 Limitations of the Work

This analysis of the current AI education domain has enabled insight into a number of difficulties both students and lectures encounter within this field. As with any other research project there are limitations, many of the limitations of this study centre around low response rate. This may potentially impact the solidity from which statements can be made addressing the research questions, particularly RQ1 and RQ2 if a large enough sample size has not been achieved. However, methodological triangulation and the diversity of roles of the participants allow for initiation of the identification of best practice and the perceived difficulties.

As participation in all aspects of this study were optional, with lecturers acting as gatekeeper when requiring student participation, there were constraints in place relating to the communication of the research to students. This included how much time they were given to complete the data collection methods such as the questionnaires. Requests were also made to analyse the module results to determine whether there were any differences in attainment for example between students who had a stronger mathematical background. However, this data was unfortunately inaccessible for the purpose of this study. Data relating to module results would have been beneficial to have some measurement relating to the effectiveness of the mitigation strategies (RQ3). However, findings from the user reviews of MetaLearning

addressed the perceived usefulness as identified by students, lecturers and computing professionals.

One of the limitations relating to data analysis was the lack of reliability assessment undertaken for all thematic analysis. Inter-rater reliability was completed for interviews with lecturers (Section 4.4.2) and for the UKICER workshop findings (Section 6.6.2). However, it could have been completed for all analysis tasks requiring thematic coding to provide greater validity to the findings from this specific method of data collection. However, when dealing with qualitative data there is “no single canon of validity” therefore “fitness for purpose within an ethically defensible framework should be adopted” (Cohen, Manion and Morrison, 2018, p.526). Rigorous research design was implemented when constructing the framework for the qualitative data collection, with a clear connection between the theory, research objectives and structure and line of questioning for the interviews.

Overall, within this study there was a proliferation of similar types of universities, both within the case studies and the online systematic review of modules as only a snapshot of universities were analysed. Courses which covered some form of AI might have been missed out for inclusion within this study as access to all of the course materials was not available online and the lecturer may be teaching aspects of this domain without explicitly stating this in the syllabus. AI and in particular, Machine Learning is commonly being taught on subjects outside of the computing domain, therefore analysis of these courses will also be beneficial to fully comprehend teaching and learning practices pertinent to this field.

8.5 Future Work

Due to the lack of current research pertaining to education for Machine Learning, there are various directions in which this work can continue and expand upon the findings outlined in this thesis. Repetition of the online review of modules outlined in Chapter 4 (Section 4.3, p.82) with a wider scope of universities and over a much broader geographical area will be beneficial to determine if the findings of this research are similar for differing cohorts of students. Further iteration of the online review will also identify other module leaders/lecturers who may be willing to participate in a case study to determine to what extent the findings of this research generalise and to determine any other barriers to learning which have not yet been

identified. Further iteration of the data collection methods of the case studies will also contribute to the solidification of the PCK and threshold concepts within this domain.

Alongside the iteration of the online review of modules, greater examination of the educational provision relating to data ethics will also be revisited due to the finding that only 10% of the Machine Learning modules analysed covered some form of ethics (Section 4.3.2, p.88). The lack of comprehension of the ethical considerations of this domain as highlighted in the UK Data Skills Gap report (Department for Digital Media and Sport, 2021) highlights the need for a potential resource which lecturers can use to guide and assist them into incorporating this topic within their modules.

The idea and importance of a clear mental model, which is the individual's mental representation (Storey, Fracchia and Müller, 1999) of, in this particular instance the AI domain has been a theme throughout the thesis, first discussed in Chapter 2 relating to the theories of learning (Section 2.5.2, p.29) and the potential preconceptions a learner may have of the domain (Section 2.6, p.38). Students participating in the case study at Newcastle University were asked what they thought Machine Learning was prior to starting the module (Section 4.5.1, p.106) and the majority had an accurate comprehension of the domain. However, one of the lecturers interviewed for the review of MetaLearning discussed how a number of students dropped out of the module when they realise the amount of mathematics involved. Further work in understanding students' comprehension of the field of AI prior to commencing any form of module or course within this area will enable greater understanding of any preconceptions surrounding the field. The findings from this future work may also be relevant to the aim of democratising and widening participation within this field as individuals may not have a thorough understanding of the breadth of the domain.

Further investigation of the learning strategies students employ when learning this domain, and within Computer Science in general will determine which methods learner perceive to be useful and relevant to their learning. This extension to the research is motivated by the findings from the case studies (Section 4.5.1, 4.5.2) which identified that the postgraduate students did not use more holistic strategies such as reflection and goal setting. However, these strategies are pertinent to aiding metacognition and becoming a more self-regulated learner.

Implementation of the feedback pertaining to MetaLearning will enable improvements to be made to the current system to assist learners in identifying areas they require further study, for example through improvement of the current feedback provision. Nicol and MacFarlane-Dick's (2006) seven principles of good practice feedback which support self-regulation will be used to improve the current feedback provision within the online tool. A tutorial which has a specific focus on how to program a selection of the Machine Learning and Deep Learning models discussed within the other tutorials will also be created influenced by the suggestions from the lecturers reviewing the online tool (Section 6.4.3, p.199). Inclusion of the programming tutorial may assist the learners in contextualising the theoretical knowledge as well as assisting students in learning the practical aspects of this domain which both cohorts of students from the case studies identified as most difficult.

The approach utilised within this research could also be applied to other Computer Science domains, for example Cryptography due to the similarity in requirements of high mathematical and technical content. Comparable to AI, there is a lack of pedagogical research relating to the best practice for teaching this particular strand of Computer Science, therefore the methodology used within this study could be applied within Cryptography or generalised to other educational disciplines to determine the pedagogical content knowledge and specific best practice.

8.6 Conclusion

To address the current lack of research pertaining to the best practices for teaching the often complex domain of AI, a range of qualitative data collection methods have been used to identify and assess current provision and pedagogical practices used. The work conducted in this thesis has helped identify barriers students face when studying this domain including mathematics anxiety and low confidence in technical skills. Pertinent to this identified barrier were the differences in cohort educational background and the need to include some form of teaching of foundational knowledge, this was a motivating factor for MetaLearning. The importance of an accurate mental model was a key barrier identified throughout this thesis, from the literature review in Chapter 2 (p.29), to the responses from the students in the case studies in Chapter 4 (p.106) and finally in comments made by lecturers reviewing the online tool in Chapter 6 (p.197). As discussed in Chapter 2 (p.38), modules within the AI domain are

primarily taught within Computer Science programmes where learners will bring a prior understanding of for example traditional programming where they have to program the steps detailing how to achieve the outcome. However, programming Machine Learning is very different from this in that the programmer defines the objective “that the system is trying to maximise” (Shapiro, Fiebrink and Norvig, 2018). Therefore, students will need to re-evaluate their mental model which may cause difficulties in comprehension of the new domain. To assist learners in overcoming these identified barriers, strategies were implemented within the online tool to build the learners metacognition and self-efficacy. Students and lecturers who reviewed MetaLearning, as detailed in Chapter 6 (p.207) felt that the knowledge surveys and the use of testing as a learning tool were useful and effective methods.

Within Chapter 7 (p.219), an initial outline of the threshold concepts have been defined. These concepts were identified through varying data collection and analysis methods outlined in Chapters 3, 4, 5 and 6. Although these findings are preliminary and will require iteration of the case studies and interviews with lecturers within this domain they provide an insight into specific topics and aspects of the domain that students struggle with. These topics include specific Machine Learning algorithms including the Support Vector Machine (Cortes and Vapnik, 1995), an issue overall with the sub-domain of Deep Learning and difficulties with the foundational mathematics and statistic knowledge. Relating to the threshold concepts and the pedagogical content knowledge of the field of Machine Learning, a framework of topics was created based on the findings from Chapter 4 (p.81) to determine essential content for an introductory course for this domain. This framework formed the content for MetaLearning and may help to alleviate the issue discussed in Section 7.3.2 (p.223) of narrowing down the content for a module within this domain.

Potential best practices were also discovered and discussed within this thesis with the aim of improving student satisfaction and attainment when studying courses within this domain. Specific practices included the idea of allowing the learners to explore the concept and ideas of Machine Learning to assist them in building a stronger more accurate mental model before concerning themselves with the programming syntax. This also converges with the highlighted importance of the inclusion of real-life examples when teaching this domain to help learners conceptualise these new, often challenging concepts. As identified in the literature review in Chapter 2 (p.38) and through observation of modules within this domain in the case studies,

Machine Learning is an inherently active learning discipline. Therefore, the inclusion of group exercises and varying activities within the sessions can promote engagement and further interaction between both peers and lecturers.

Further iteration and extension of this research will help solidify the findings from this research and continue to define the barriers, threshold concepts and pedagogical content knowledge of this domain to assist and potentially improve the experiences of both lecturers and students within this domain.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A. and Irving, G. (2015) *TensorFlow: Large Scale Machine Learning on Heterogeneous Systems*, *TensorFlow*. Available at: <https://www.tensorflow.org/> (Accessed: 9 November 2021).

Abiodun, O.I., Jantan, A., Omolara, A.E., Dada, K.V., Mohamed, N.A.E. and Arshad, H. (2018) 'State-of-the-Art in Artificial Neural Network Applications: A Survey', *Helijon*, 4(11), p. e00938. doi:10.1016/j.heliyon.2018.e00938.

ACM (2018) *Fathers of the Deep Learning Revolution Receive ACM A.M. Turing Award*. Available at: <https://awards.acm.org/about/2018-turing> (Accessed: 27 July 2020).

ACM (2021) *Words Matter: Alternatives for Charged Terminology in the Computing Profession*. Available at: <https://www.acm.org/diversity-inclusion/words-matter> (Accessed: 22 September 2021).

Akinfaderin, W. (2017) *The Mathematics of Machine Learning, Towards Data Science*. Available at: <https://towardsdatascience.com/the-mathematics-of-machine-learning-894f046c568> (Accessed: 26 November 2020).

Algorithmia (2020) *What is Bayesian Machine Learning?*, *Algorithmia Blog*. Available at: <https://algorithmia.com/blog/bayesian-machine-learning> (Accessed: 14 December 2021).

Allen, B., Heels, L. and Devlin, M. (2021) 'Measuring the Difference Between Student and Staff Perception of Self-Efficacy and Confidence Using Online Tools', in *UKICER*, p. 53. doi:10.1145/792548.611930.

Altman, D.G. and Bland, J.M. (2009) 'Parametric v non-parametric methods for data analysis', *BMJ (Online)*. British Medical Journal Publishing Group, p. 170. doi:10.1136/bmj.a3167.

Armstrong, P. (2010) *Bloom's Taxonomy*, *Vanderbilt University Center for Teaching*. Available

at: <https://cft.vanderbilt.edu/guides-sub-pages/blooms-taxonomy/> (Accessed: 28 July 2021).

Artino, A. and Ioannou, A. (2008) 'Promoting Academic Motivation and Self-Regulation: Practical Guidelines for Online Instructors', *Society for information technology & teacher education international conference. Association for the Advancement of Computing in Education (AACE)*, pp. 1–11.

Asamoah, D., Doran, D. and Schiller, S. (2015) 'Teaching the Foundations of Data Science: An Interdisciplinary Approach', in *Pre-ICIS SIGDSA Workshop*. Texas, pp. 1–9. Available at: <http://arxiv.org/abs/1512.04456>.

Ashcraft, M., Kirk, E. and Hopko, D. (1998) 'On the cognitive consequences of mathematics anxiety', *The development of mathematical skills. Studies in developmental psychology.*, (November), pp. 175–196.

Baillie, C., Bowden, J.A. and Meyer, J.H.F. (2013) 'Threshold capabilities: Threshold concepts and knowledge capability linked through variation theory', *Higher Education*, 65(2), pp. 227–246. doi:10.1007/s10734-012-9540-5.

Bandura, A. (1977) 'Self-efficacy: Toward a Unifying Theory of Behavioral Change', *Psychological Review*, 84(2), pp. 191–215.

Barbosa, H. and Garcia, F. (2005) 'Importance of online assessment in the E-learning process', *ITHET 2005: 6th International Conference on Information Technology Based Higher Education and Training, 2005*, 2005, pp. 7–12. doi:10.1109/ITHET.2005.1560287.

Barnes, J., Colquhoun, J., Devlin, M., Heels, L., Lord, P., Marshall, L., Napier, C., Solaiman, E., Speirs, N., Talbot, L., Ushaw, G. and Witty, C. (2020) 'Designing a Portfolio-Oriented Curriculum using Problem Based Learning', in *Computing Education Practice 2020*. New York: ACM. doi:10.1145/3372356.3372367.

Baturay, M.H. (2015) 'An Overview of the World of MOOCs', *Procedia - Social and Behavioral Sciences*, 174, pp. 427–433. doi:10.1016/j.sbspro.2015.01.685.

Beck, K. (2018) *What Is the Tukey HSD Test?*, *Sciencing*. Available at: <https://sciencing.com/what-is-the-tukey-hsd-test-12751748.html> (Accessed: 5 January 2022).

Bengio, Y. and Ford, M. (2018) 'Yoshua Bengio', in *Architects of Intelligence*. Birmingham: Packt Publishing.

Bengio, Y. and Grandvalet, Y. (2004) 'No Unbiased Estimator of the Variance of K-Fold Cross-Validation', *Journal of Machine Learning Research*, 5, pp. 1089–1105.

Blackboard (2021) *Blackboard Learn: An Advanced LMS Blackboard*. Available at: <https://www.blackboard.com/en-uk/teaching-learning/learning-management/blackboard-learn#main-content> (Accessed: 19 September 2021).

BlackBoard Inc. (2022) *About Us*. Available at: <https://www.blackboard.com/en-uk/about-us> (Accessed: 19 May 2022).

Bock, H. (2008) 'Origins and Extensions of the K-Means Algorithm in Cluster Analysis', *Electronic Journal for History of Probability and Statistics*, 4(2), pp. 1–18.

Boekaerts, M. (1997) 'Self-regulated learning: A new concept embraced by researchers, policy makers, educators, teachers, and students', *Learning and Instruction*, 7(2), pp. 161–186. doi:10.1016/S0959-4752(96)00015-1.

Boese, E. (2016) 'Just-In-Time Learning for the Just Google It Era', *Proceedings of the 47th ACM Technical Symposium on Computing Science Education - SIGCSE '16*, pp. 341–345. doi:10.1145/2839509.2844583.

Bostock, S.J. (1998) 'Constructivism in mass higher education: a case study', *British Journal of Educational Technology*, 29(3), p. 11.

Bourne, J. (2019) *93% of firms committed to AI - but skills shortage posing problems*. Available at: <https://artificialintelligence-news.com/2019/05/21/93-of-firms-committed-to-ai-but->

skills-shortage-posing-problems/ (Accessed: 23 June 2020).

Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101. doi:10.1191/1478088706qp063oa.

Brenner, M. (2006) 'Interviewing in Educational Research', in Camilli, G., Green, J., and Elmore, P. (eds) *Handbook of Complementary Methods in Education Research*. Abingdon: Routledge.

Brownlee, J. (2016) *Crash Course On Multi-Layer Perceptron Neural Networks, Machine Learning Mastery*. Available at: <https://machinelearningmastery.com/neural-networks-crash-course/> (Accessed: 16 December 2021).

Brownlee, J. (2017) *A Gentle Introduction to Long Short-Term Memory Networks by the Experts, Machine Learning Mastery*. Available at: <https://machinelearningmastery.com/gentle-introduction-long-short-term-memory-networks-experts/> (Accessed: 14 December 2021).

Brunner, R.J. and Kim, E.J. (2016) 'Teaching data science', *Procedia Computer Science*, 80, pp. 1947–1956. doi:10.1016/j.procs.2016.05.513.

Budhiraja, A. (2016) *Dropout in (Deep) Machine learning, Medium*. Available at: <https://medium.com/@amarbudhiraja/https-medium-com-amarbudhiraja-learning-less-to-learn-better-dropout-in-deep-machine-learning-74334da4bfc5> (Accessed: 12 November 2021).

Cameron, D. and Maguire, K. (2017) *Public views of Machine Learning: Digital Natives, Royal Society*. Available at: <https://royalsociety.org/topics-policy/projects/machine-learning/> (Accessed: 6 January 2019).

Canas, J., Bajo, M. and Gonzalvo, P. (1994) 'Mental models and computer programming', *Human Computer Studies*, 40, pp. 795–811. doi:10.1063/1.3033202.

Centre for Data Ethics and Innovation (2020) *AI Barometer Report*, GOV.UK. Available at: <https://www.gov.uk/government/publications/cdei-ai-barometer> (Accessed: 2 February

2022).

Cera, R., Mancini, M. and Antonietti, A. (2014) 'Relationships between Metacognition, Self-efficacy and Self-regulation in Learning', *ECPS - Educational, Cultural and Psychological Studies*, (7), pp. 115–141. doi:10.7358/ecps-2013-007-cera.

Ceron, R. (2019) *AI, Machine Learning and Deep Learning: What's the Difference?*, IBM. Available at: <https://www.ibm.com/blogs/systems/ai-machine-learning-and-deep-learning-whats-the-difference/> (Accessed: 3 December 2020).

Chizmar, J.F. and Ostrosky, A.L. (1998) 'The One-Minute Paper: Some Empirical Findings', *Journal of Economic Education*, 29(1), pp. 3–10. doi:10.1080/00220489809596436.

Chollet, F. (2015) *Keras*, *Github*. Available at: <https://github.com/keras-team/keras> (Accessed: 26 November 2020).

Chollet, F. (2018) *Deep Learning with Python*. New York: Manning.

Clow, D. (2013) 'MOOCs and the Funnel of Participation', in *Proceedings of the third international conference on learning analytics and knowledge*. Leuven, pp. 185–189.

Code.org (2021a) *AI and Machine Learning Module*. Available at: <https://studio.code.org/s/aiml-2021> (Accessed: 3 August 2021).

Code.org (2021b) *Learn today, build a brighter tomorrow*. Available at: <https://code.org/> (Accessed: 3 August 2021).

Cohen, L., Manion, L. and Morrison, K. (2018) *Research Methods in Education*. 8th edn. New York: Routledge.

Complete University Guide (2021) *Computer Science - Top UK University Subject Tables and Rankings 2021*. Available at: <https://www.thecompleteuniversityguide.co.uk/league-tables/rankings/computer-science> (Accessed: 23 February 2021).

Concannon, F., Flynn, A. and Campbell, M. (2005) 'What campus-based students think about the quality and benefits of e-learning', *British Journal of Educational Technology*, 36(3), pp. 501–512.

Cortes, C. and Vapnik, V. (1995) 'Support-Vector Networks', *Machine Learning*, 20, pp. 273–297. doi:<https://doi.org/10.1007/BF00994018>.

Coursera (2021) *Coursera - Build Skills with Online Courses from Top Institutions*. Available at: <https://www.coursera.org/> (Accessed: 4 August 2021).

Coursera (2022) *Introduction to Artificial Intelligence (AI)*. Available at: <https://www.coursera.org/learn/introduction-to-ai?action=enroll> (Accessed: 22 May 2022).

Cousin, G. (2006) 'An Introduction to Threshold Concepts', *Planet*, 17(1), pp. 4–5. doi:[10.11120/plan.2006.00170004](https://doi.org/10.11120/plan.2006.00170004).

Cramer, J. (2002) 'The Origins of Logistic Regression', *Tinbergen Institute Discussion Paper*, 119(4). Available at: <https://papers.tinbergen.nl/02119.pdf>.

Crawford, J., Butler-Henderson, K., Jurgen, R., Malkawi, B., Glowatz, M., Burton, R., Magni, A.P. and Lam, S. (2020) 'COVID-19 : 20 countries ' higher education intra-period digital pedagogy responses', *Journal of Applied Learning & Teaching*, 3(1).

Crick, T. (2017) 'Computing Education: An Overview of Research in the Field', *Cronfa - Swansea University Open Access Repository*, (April), pp. 1–38. Available at: <http://cronfa.swan.ac.uk/Record/cronfa43589%0Ahttps://royalsociety.org/-/media/policy/projects/computing-education/literature-review-overview-research-field.pdf>.

Curry, N. (2018) *Putting the pedagogy first in digital pedagogies*, Cambridge University Press. Available at: <https://www.cambridge.org/elt/blog/2018/10/05/putting-the-pedagogy-first-in-digital-pedagogies/> (Accessed: 16 October 2020).

Curzon, P., McOwan, P., Donohue, J., Wright, S. and Marsh, W. (2018) 'Teaching of Concepts', in Sentence, S., Barendsen, E., and Schulte, C. (eds) *Computer Science Education- Perspectives*

on *Teaching and Learning in School*. London: Bloomsbury Academic.

D'Ignazio, C. and Klein, L.F. (2020) *Data Feminism*. Massachusetts: The MIT Press.

David, M. (2010) 'Introduction to the dilemmas of widening participation in higher education', in David, M. (ed.) *Improving Learning by Widening Participation in Higher Education*. New York: Routledge.

Davies, P. (2006) 'Threshold Concepts: How Can We Recognise Them?', in Land, R. and Meyer, J.H.F. (eds) *Overcoming Barriers to Student Understanding*. Abingdon: Routledge.

Deisenroth, M.P., Faisal, A.A. and Ong, C.S. (2020) *Mathematics For Machine Learning*. Cambridge: Cambridge University Press.

Deloitte and Institute of Coding (2018) *Diversifying Digital Learning*. doi:10.1353/book.57781.

Demetriou, A. and Spanoudis, G. (2018) *Growing Minds - A Developmental Theory of Intelligence, Brain, and Education*. New York: Routledge.

Department for Digital Media and Sport (2017) *UK Digital Strategy*, Government of the United Kingdom. Available at: <https://www.gov.uk/government/publications/uk-digital-strategy/uk-digital-strategy> (Accessed: 21 September 2020).

Department for Digital Media and Sport (2021) *Quantifying the UK Data Skills Gap*, GOV.UK. Available at: <https://www.gov.uk/government/publications/quantifying-the-uk-data-skills-gap/quantifying-the-uk-data-skills-gap-full-report> (Accessed: 21 May 2021).

Department for Education (2014) *Mathematics Programmes of Study: Key Stage 4 National Curriculum in England*. Available at: <https://www.gov.uk/government/publications/national-curriculum-in-england-mathematics-programmes-of-study> (Accessed: 21 September 2021).

Discord (2021) *Discord - Your Place to Talk and Hang Out*. Available at: <https://discord.com/> (Accessed: 21 September 2021).

Dixon, L. and O’Gorman, V. (2020) “‘Block teaching’—exploring lecturers’ perceptions of intensive modes of delivery in the context of undergraduate education’, *Journal of Further and Higher Education*, 44(5), pp. 583–595. doi:10.1080/0309877X.2018.1564024.

Doo, M.Y., Bonk, C. and Heo, H. (2020) ‘A Meta-Analysis of Scaffolding Effects in Online Learning in Higher Education’, *International Review of Research in Open and Distributed Learning*, 21(3), pp. 60–80.

Downing, K., Ho, R., Shin, K., Vrijmoed, L. and Wong, E. (2007) ‘Metacognitive development and moving away’, *Educational Studies*, 33(1), pp. 1–13. doi:10.1080/03055690600850347.

Draper, S. and Waldman, J. (2013) *Deep and Surface Learning: The Literature*, CERE12-13. Available at: <https://www.psy.gla.ac.uk/~steve/courses/archive/CERE12-13-safari-archive/topic9/webarchive-index.html> (Accessed: 28 July 2021).

Drummond, S.A. (2009) *Investigating the Impact of Entry Qualifications on Student Performance in Computing Programmes at Undergraduate Level*. Durham University. Available at: <http://etheses.dur.ac.uk/154/>.

Duke, B., Harper, G. and Johnston, M. (2013) ‘Connectivism as a Digital Age Learning Theory’, *The International HETL Review*, (Special Issue), pp. 4–13. doi:10.1016/0007-6813(92)90075-K. Durham University (2021) *Faculty Handbook Archive : Archive Module Description*. Available at:

https://www.dur.ac.uk/faculty.handbook/archive/module_description/?year=2020&module_code=COMP3547 (Accessed: 9 November 2021).

Dweck, C.S. (2017) *Mindset*. Revised. London: Robinson.

Eckerdal, A., Kinnunen, P., Thota, N., Nylén, A., Sheard, J. and Malmi, L. (2014) ‘Teaching and learning with MOOCs: Computing academics’ perspectives and engagement’, in *ITICSE 2014 - Proceedings of the 2014 Innovation and Technology in Computer Science Education Conference*, pp. 9–14. doi:10.1145/2591708.2591740.

Eddy, S. (2010) *Effect size for Analysis of Variance (ANOVA)*, *Psycho Hawks*. Available at: <https://psychohawks.wordpress.com/2010/10/31/effect-size-for-analysis-of-variables-anova/> (Accessed: 23 January 2022).

Editors of Encyclopedia Britannica (2013) *Higher Education*, *Britannica*. Available at: <https://www.britannica.com/topic/higher-education> (Accessed: 11 November 2021).

edX (2021) *edX - Free Online Courses by Harvard, MIT, & more*. Available at: <https://www.edx.org/> (Accessed: 4 August 2021).

Encyclopedia Britannica (2021) *Turing test / Definition & Facts / Britannica*. Available at: <https://www.britannica.com/technology/Turing-test> (Accessed: 29 May 2020).

Engeström, Y. (2018) 'Expansive Learning - Towards an activity-theoretical reconceptualization', in Illeris, K. (ed.) *Contemporary Theories of Learning*. New York: Routledge.

Erickson, S. and Heit, E. (2015) 'Metacognition and confidence: Comparing math to other academic subjects', *Frontiers in Psychology*, 6(JUN), pp. 1–10. doi:10.3389/fpsyg.2015.00742.

European Commission (2020) *On Artificial Intelligence - A European approach to excellence and trust*. Brussels. Available at: https://ec.europa.eu/info/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en (Accessed: 2 February 2022).

Eynon, R. and Young, E. (2020) 'Methodology, Legend, and Rhetoric: The Constructions of AI by Academia, Industry, and Policy Groups for Lifelong Learning', *Science Technology and Human Values*, pp. 1–26. doi:10.1177/0162243920906475.

Fahlman, S.E., Hinton, G.E. and Sejnowski, T.J. (1983) 'Massively Parallel Architectures for AI: NETL, Thistle, and Boltzmann Machines.', *AAAI-83 Proceedings*, pp. 109–113.

Fayyad, U. and Hamutcu, H. (2020) 'Toward Foundations for Data Science and Analytics: A Knowledge Framework for Professional Standards', *Harvard Data Science Review*, 2(2).

doi:<https://doi.org/10.1162/99608f92.1a99e67a>.

Feigenbaum, E.A. (1963) 'Artificial Intelligence Research', *IEEE Transactions of the Professional Technical Group on Information Theory*, 1(4).

Fiebrink, R. (2019) 'Machine learning education for artists, musicians, and other creative practitioners', *ACM Transactions on Computing Education*, 19(4). doi:10.1145/3294008.

Fink, A. (2010) 'Survey Research Methods', *International Encyclopedia of Education*, pp. 152–160. doi:10.1016/B978-0-08-044894-7.00296-7.

Forsey, M., Low, M. and Glance, D. (2013) 'Flipping the sociology classroom: Towards a practice of online pedagogy', *Journal of Sociology*, 49(4), pp. 471–485. doi:10.1177/1440783313504059.

Foster, B., Perfect, C. and Youd, A. (2012) 'A completely client-side approach to e-assessment and e-learning of mathematics and statistics', *International Journal of e-Assessment*, 2(2). Available at: <http://journals.sfu.ca/ijea/index.php/journal/article/view/35>.

Freeman, J. and Campbell, M. (2007) 'The Analysis of Categorical Data : Fisher's Exact Test', *Scope*, 33(5), pp. 11–12. Available at: http://www.sheffield.ac.uk/polopoly%7B_%7Dfs/1.43998!/file/tutorial-9-fishers.pdf http://www.sheffield.ac.uk/polopoly_fs/1.43998!/file/tutorial-9-fishers.pdf.

Freire, P. (1970) *Pedagogy of the Oppressed*. New York: Continuum Publishing Company.

Freire, P. (2004) 'The "Banking" Concept of Education', in Canestrari, A.S. and Marlowe, B.A. (eds) *Educational Foundations: An Anthology of Critical Readings*. London: Sage, pp. 99–111.

Fujita, H. (2020) 'AI-based computer-aided diagnosis (AI-CAD): the latest review to read first', *Radiological Physics and Technology*, 13(1), pp. 6–19. doi:10.1007/s12194-019-00552-4.

Gandhi, R. (2018) *Naive Bayes Classifier, Towards Data Science*. Available at:

<https://towardsdatascience.com/naive-bayes-classifier-81d512f50a7c> (Accessed: 14 December 2021).

Geert van den Berg, R. (2021) *Effect Size in Statistics, SPSS Tutorials*. Available at: <https://www.spss-tutorials.com/effect-size/> (Accessed: 25 January 2022).

Géron, A. (2017) *Hands-On Machine Learning with Scikit-Learn & TensorFlow*. 1st edn. Edited by N. Tache. California: O'Reilly.

Gleason, N. (2018a) *Higher Education In The Era Of The Fourth Industrial Revolution*. Edited by N. Gleason, W. Singapore: Palgrave Macmillan.

Gleason, N. (2018b) 'Singapore's Higher Education Systems in the Era of the Fourth Industrial Revolution: Preparing Lifelong Learners', in Gleason. W, N. (ed.) *Higher Education In the Era Of The Fourth Industrial Revolution*. Singapore: Palgrave Macmillan.

Gluga, R., Kay, J., Lister, R., Simon and Kleitman, S. (2013) 'Mastering cognitive development theory in computer science education', *Computer Science Education*, 23(1), pp. 24–57. doi:10.1080/08993408.2013.768830.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. (2020) 'Generative Adversarial Networks', *Communications of the ACM*, 63(11), pp. 139–144. doi:10.1145/3422622.

Goodfellow, I.J., Bengio, Y. and Courville, A. (2016) *Deep Learning*. Massachusetts: The MIT Press.

Google (2021) *Teachable Machine*. Available at: <https://teachablemachine.withgoogle.com/> (Accessed: 3 August 2021).

Greca, I.M. and Moreira, M.A. (2000) 'Mental models, conceptual models, and modelling', *International Journal of Science Education*, 22(1), p. 11. doi:10.1080/095006900289976.

Grow, G.O. (1991) 'Teaching learners to be self-directed', *Adult Education Quarterly*, 41(3), pp. 125–149. doi:10.1177/0001848191041003001.

Guzdial, M. (2021) 'Talking about race in CS education', *Communications of the ACM*, 64(1), pp. 10–11. doi:10.1145/3433921.

Hall, W. and Pesenti, J. (2017) *Growing the Artificial Intelligence Industry in the UK*. Available at:

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf <https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in->.

Hardin, J., Hoerl, R., Horton, N.J., Nolan, D., Baumer, B., Hall-Holt, O., Murrell, P., Peng, R., Roback, P., Temple Lang, D. and Ward, M.D. (2015) 'Data Science in Statistics Curricula: Preparing Students to "Think with Data"', *American Statistician*, 69(4), pp. 343–353. doi:10.1080/00031305.2015.1077729.

Hattie, J. (2012) *Visible Learning For Teachers*. Abingdon: Routledge.

Hattie, J. and Donoghue, G. (2018) 'A Model of Learning - Optimizing the effectiveness of learning strategies', in Illeris, K. (ed.) *Contemporary Theories of Learning*. New York: Routledge.

Hattie, J.A.C. and Donoghue, G.M. (2016) 'Learning strategies: a synthesis and conceptual model', *Nature Partner Journals Science of Learning*, 1(1). doi:10.1038/npjscilearn.2016.13.

Hazzan, O., Lapidot, T. and Ragonis, N. (2011) *Guide to Teaching Computer Science: An Activity Based Approach*. New York: Springer. doi:10.1007/978-1-4471-6630-6.

Hembree, R. (1990) 'The Nature, Effects, and Relief of Mathematics Anxiety', *Journal for Research in Mathematics Education*, 21(1), pp. 33–46.

Henderson, T. (2019) 'Teaching Data Ethics', in *Computing Education Practice*, pp. 1–4.

HESA (2019) *Higher Education Student Statistics: UK, 2018/19 - Subjects studied*, HESA. Available at: <https://www.hesa.ac.uk/news/16-01-2020/sb255-higher-education-student-statistics/subjects> (Accessed: 27 November 2021).

HESA (2021) *What do HE students study?*, HESA. Available at: <https://www.hesa.ac.uk/data-and-analysis/students/what-study> (Accessed: 27 November 2021).

Hicks, S.C. and Irizarry, R.A. (2017) 'A Guide to Teaching Data Science', *The American Statistician*, 1305, pp. 1–34. doi:10.1080/00031305.2017.1356747.

HM Government (2018) *Industrial Strategy Artificial Intelligence Sector Deal*. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/702810/180425_BEIS_AI_Sector_Deal__4_.pdf (Accessed: 21 September 2021).

Ho, T.K. (1995) 'Random Decision Forests', *Proceedings of 3rd International Conference on Document Analysis and Recognition*, pp. 278–282. doi:10.1109/ICDAR.1995.598994.

Hoban, S. and Hoban, G. (2004) 'Self-Esteem, Self-Efficacy and Self-Directed Learning: Attempting to Undo the Confusion', *International Journal of Self-Directed Learning*, 1(2), pp. 7–25.

Hochreiter, S. and Urgen Schmidhuber, J. (1997) 'Long Short-Term Memory', *Neural Computation*, 9(8), pp. 1735–1780. doi:10.1162/neco.1997.9.8.1735.

Hsu, F., Campbell, Murray, S. and Hoane, A.J. (1995) 'Deep Blue System Overview', in *Proceedings of the 9th International Conference on Supercomputing*, pp. 240–244.

Huang, J., Li, Y.F. and Xie, M. (2015) 'An empirical analysis of data preprocessing for machine learning-based software cost estimation', *Information and Software Technology*, 67, pp. 108–127. doi:10.1016/j.infsof.2015.07.004.

Hubbard, A. (2018) 'Pedagogical content knowledge in computing education: a review of the research literature', *Computer Science Education*, 28(2), pp. 117–135. doi:10.1080/08993408.2018.1509580.

Hunt, T., Clark-Carter, D. and Sheffield, D. (2011) 'The Development and Part Validation of a U.K. Scale for Mathematics Anxiety', *Journal of Psychoeducational Assessment*, 29(5), pp. 455–466. doi:<https://doi.org/10.1177/0734282910392892>.

IBM (2020a) *Closing the AI skills gap*, *Financial Times*. Available at: <https://www.ft.com/partnercontent/ibm/closing-the-ai-skills-gap.html> (Accessed: 29 October 2021).

IBM (2020b) *Deep Blue*. Available at: <https://www.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/transform/> (Accessed: 18 June 2020).

IBM Corp (2019) 'IBM SPSS Statistics for Mac, Version 26.0'. Armonk, New York: IBM Corp.

Illeris, K. (2007) *How We Learn: Learning and Non-Learning in School and Beyond*. London: Routledge.

Illeris, K. (2018) 'A Comprehensive Understanding of Human Learning', in Illeris, K. (ed.) *Contemporary Theories of Learning*. 2nd Ed. New York: Routledge.

Instructure (2021) *Canvas Overview*. Available at: <https://www.instructure.com/en-gb/canvas> (Accessed: 19 September 2021).

Instructure (2022) *Canvas For Higher Education*. Available at: <https://www.instructure.com/higher-education/products/canvas/canvas-lms> (Accessed: 19 May 2022).

Lossi, L. (2007) 'Strategies for Reducing Math Anxiety in Post-Secondary Students', *Proceedings of the Sixth Annual College of Education Research Conference: Urban and International*

Education Section, pp. 30–35. Available at:
http://coeweb.fiu.edu/research_conference/%5Cnhttp://digitalcommons.fiu.edu/cgi/viewcontent.cgi?article=1257&context=sferc.

Ipsos MORI (2017) 'Public views of Machine Learning. Findings from public research and engagement conducted on behalf of the Royal Society', (April), p. 87. Available at: <http://www.ipsos-mori.com/terms>.

Ironsmith, M., Marva, J., Harju, B. and Eppler, M. (2003) 'Motivation and Performance in College Students Enrolled in Self-Paced Versus Lecture-Format Remedial Mathematics Courses', *Journal of Instructional Psychology*, 30(4), pp. 276–284. doi:10.1016/j.jaci.2012.05.050.

Ivanova, M., Groszeck, G. and Holotescu, C. (2015) 'Researching Data Privacy Models in eLearning', *2015 International Conference on Information Technology Based Higher Education and Training, ITHET 2015* [Preprint]. doi:10.1109/ITHET.2015.7218033.

Jain, S. and Dowson, M. (2009) 'Mathematics anxiety as a function of multidimensional self-regulation and self-efficacy', *Contemporary Educational Psychology*, 34(3), pp. 240–249. doi:10.1016/j.cedpsych.2009.05.004.

Jain, V. (2019) *Everything you need to know about "Activation Functions" in Deep learning models, Towards Data Science*. Available at: <https://towardsdatascience.com/everything-you-need-to-know-about-activation-functions-in-deep-learning-models-84ba9f82c253> (Accessed: 12 November 2021).

Jeffcock, P. (2018) *What's the Difference Between AI, Machine Learning, and Deep Learning?*, Oracle. Available at: <https://blogs.oracle.com/bigdata/difference-ai-machine-learning-deep-learning> (Accessed: 3 December 2020).

Jolliffe, I.T. and Cadima, J. (2016) 'Principal Component Analysis: A Review and Recent Developments', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065). doi:10.1098/rsta.2015.0202.

Jones, P. (2010) *Introducing Neuroeducational Research*. Abingdon: Routledge.

Jones, R. and Thomas, L. (2005) 'The 2003 UK Government Higher Education White Paper: A critical assessment of its implications for the access and widening participation agenda', *Journal of Education Policy*, 20(5), pp. 615–630. doi:10.1080/02680930500222477.

Jordan, K. (2014) 'Initial trends in enrolment and completion of massive open online courses', *International Review of Research in Open and Distance Learning*, 15(1), pp. 133–160. doi:10.19173/irrodl.v15i1.1651.

Juhler, S.M., Rech, J.F., From, S.G. and Brogan, M.M. (1998) 'The effect of optional retesting on college students' achievement in an individualized algebra course', *Journal of Experimental Education*, 66(2), pp. 125–137. doi:10.1080/00220979809601399.

Kaper, H., Weibel, E. and Tipei, S. (1999) 'Data Sonification and Sound Visualization', *Computing in Science and Engineering*, 1(4), pp. 48–58.

Katoch, S., Chauhan, S.S. and Kumar, V. (2020) 'A Review on Genetic Algorithm: Past, Present, and Future', *Multimedia Tools and Applications*, 80, pp. 8091–8126. doi:10.1007/s11042-020-10139-6.

Kegan, R. (2018) 'What "Form" Transforms? A Constructive-Developmental Approach to Transformative Learning', in Illeris, K. (ed.) *Contemporary Theories of Learning*. New York: Routledge.

Kiley, M. and Wisker, G. (2009) 'Threshold concepts in research education and evidence of threshold crossing', *Higher Education Research and Development*, 28(4), pp. 431–441. doi:10.1080/07294360903067930.

Kirschner, P. and Hendrick, C. (2020) *How Learning Happens: Seminal Works in Educational Psychology and What They Mean in Practice*. London: Routledge.

Knowles, M.S., Holton, F. E. and Swanson, A. R. (2014) 'Introduction to Adult Learning', in *The*

Adult Learner: The Definitive Classic in Adult Education and Human Resource Development.
8th edn. New York: Routledge.

Kolb, D.A. (2015) *Experiential Learning: Experience as the Source of Learning and Development*. 2nd edn. New Jersey: Pearson Education.

Koppelman, H. (2008) 'Pedagogical Content Knowledge and Educational Cases in Computer Science: An Exploration', *Proceedings of the 2008 InSITE Conference* [Preprint]. doi:10.28945/3228.

Kostadinov, S. (2019) *Understanding Backpropagation Algorithm, Towards Data Science*. Available at: <https://towardsdatascience.com/understanding-backpropagation-algorithm-7bb3aa2f95fd> (Accessed: 12 November 2021).

KPMG and National Cyber Security Centre (2020) *Decrypting Diversity- Diversity and Inclusion in Cyber Security*. Available at: <https://home.kpmg/uk/en/home/insights/2020/06/diversity-and-inclusion-in-cyber-security-report-2020.html>.

Krathwohl, D.R. (2002) 'A Revision of Bloom's Taxonomy: An otherview', *Theory Into Practice*, 41(4), pp. 212–218. doi:10.1207/s15430421tip4104.

Kwiatkowski, R. (2021) *Gradient Descent Algorithm: A Deep Dive, Towards Data Science*. Available at: <https://towardsdatascience.com/gradient-descent-algorithm-a-deep-dive-cf04e8115f21> (Accessed: 4 May 2022).

Land, R., Cousin, G., Meyer, J.H.F. and Davies, P. (2005) 'Threshold Concepts and troublesome knowledge(3): implications for course design and evaluation', *Improving Student Learning and Inclusivity* [Preprint]. doi:10.4324/9780203966273.

Lavy, S. (2017) 'Who benefits from group work in higher education? An attachment theory perspective', *Higher Education*, 73(2), pp. 175–187. doi:10.1007/s10734-016-0006-z.

LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. (1998) 'Gradient-based learning applied to

document recognition', *Proceedings of the IEEE*, 86(11), pp. 2278–2323. doi:10.1109/5.726791.

Lecun, Y., Bengio, Y. and Hinton, G. (2015) 'Deep learning', *Nature*, 521(7553), pp. 436–444. doi:10.1038/nature14539.

Lee, J. (2009) 'Universals and specifics of math self-concept, math self-efficacy, and math anxiety across 41 PISA 2003 participating countries', *Learning and Individual Differences*, 19(3), pp. 355–365. doi:10.1016/j.lindif.2008.10.009.

Lewis, P. (2018) 'Globalizing the Liberal Arts: Twenty-First Century Education', in Gleason. W, N. (ed.) *Higher Education In the Era Of The Fourth Industrial Revolution*. Singapore: Palgrave Macmillan.

Little, J.L. and Bjork, E.L. (2015) 'Optimizing multiple-choice tests as tools for learning', *Memory and Cognition*, 43(1), pp. 14–26. doi:10.3758/s13421-014-0452-8.

Loo, C.W. and Choy, J.L.F. (2013) 'Sources of Self-Efficacy Influencing Academic Performance of Engineering Students', *American Journal of Educational Research*, 1(3), pp. 86–92. doi:10.12691/education-1-3-4.

López-Pérez, M.V., Pérez-López, M.C. and Rodríguez-Ariza, L. (2011) 'Blended learning in higher education: Students' perceptions and their relation to outcomes', *Computers and Education*, 56(3), pp. 818–826. doi:10.1016/j.compedu.2010.10.023.

Loui, M. and Borrego, M. (2019) 'Engineering Education Research', in Fincher, S.A. and Robins, A. (eds) *The Cambridge Handbook of Computing Education Research*. Cambridge: Cambridge University Press, pp. 292–321.

Lucas, U. and Mladenovic, R. (2007) 'The potential of threshold concepts: an emerging framework for educational research and practice', *London Review of Education*, 5(3), pp. 237–248. doi:10.1080/14748460701661294.

Luckin, R. (2018) *Machine Learning and Human Intelligence*. London: UCL Institute of Education Press.

Manikandan, S. (2011) 'Measures of Central Tendency: Median and Mode', *Journal of Pharmacology & Pharmacotherapeutics*, 2(3), pp. 214–215. doi:10.4103/0976-500x.83300.

Marshall, E. (2021) ANOVA in SPSS, Stats Tutor. Available at: https://www.sheffield.ac.uk/polopoly_fs/1.531211!/file/MASH_Oneway_ANOVA_SPSS.pdf (Accessed: 8 February 2021).

Martínez-Plumed, F., Loe, B.S., Flach, P., Eéigeartaigh, S., Vold, K. and Hernández-Orallo, J. (2018) 'The facets of artificial intelligence: A framework to track the evolution of AI', *IJCAI International Joint Conference on Artificial Intelligence*, 2018-July, pp. 5180–5187. doi:10.24963/ijcai.2018/718.

Matfield, K. (2016) *Gender Decoder: find subtle bias in job ads*. Available at: <http://gender-decoder.katmatfield.com/> (Accessed: 22 September 2021).

Maynard, A.D. (2015) 'Navigating the fourth industrial revolution', *Nature Nanotechnology*, 10(12), pp. 1005–1006. doi:10.1038/nnano.2015.286.

McCarthy, J. (1959) 'Programs with Common Sense', *Stanford University*, pp. 1–15. doi:<http://www-formal.stanford.edu/jmc/>.

McCarthy, J., Minsky, M.L., Rochester, N. and Shannon, C.E. (2006) 'A proposal for the Dartmouth summer research project on artificial intelligence', *AI Magazine*, 27(4), pp. 12–14. McCarthy, M. (2016) 'Experiential Learning Theory: From Theory To Practice', *Journal of Business & Economics Research (JBER)*, 14(3), pp. 91–100. doi:10.19030/jber.v14i3.9749.

McHugh, M. (2012) 'Interrater Reliability: The Kappa Statistic', *Biochemia Medica*, 22(3), pp. 276–282.

McLean, A. (2009) *Motivating Every Learner*. London: SAGE Publications.

Meyer, J.H.F. and Land, R. (2003) 'Threshold concepts and troublesome knowledge: Linkages to ways of thinking and practising within the disciplines', *Improving Student Learning: Improving Student Learning Theory and Practice—Ten Years On.*, pp. 412–424. doi:10.1007/978-3-8348-9837-1.

Microsoft (2021) *Microsoft Teams*. Available at: <https://www.microsoft.com/en-gb/microsoft-teams/group-chat-software> (Accessed: 11 November 2021).

Miller, Craig, S., Lehman, Jill, F. and Koedinger, Kenneth, R. (1999) 'Goals and learning in microworlds', *Cognitive Science*, 23(3), pp. 305–336. doi:10.1016/s0364-0213(99)00007-5.

Minitab (2019) *A Comparison of the Pearson and Spearman Correlation Methods*. Available at: <https://support.minitab.com/en-us/minitab-express/1/help-and-how-to/modeling-statistics/regression/supporting-topics/basics/a-comparison-of-the-pearsong-and-spearman-correlation-methods/> (Accessed: 3 April 2021).

Minsky, M., Singh, P. and Sloman, A. (2004) 'The St. Thomas Common Sense Symposium: Designing Architectures for Human-Level Intelligence', *AI Magazine*, 25(2), pp. 113–124.

MIT (2021) *Machine Learning at MIT - Classes*. Available at: <http://machinelearning.mit.edu/classes2.html> (Accessed: 9 November 2021).

Molenda, M., Reigeluth, C. and Nelson, L.M. (1983) 'Instructional Design', *Instructional Design*, 1, pp. 574–578.

Moodle (2021) *Moodle - Open-source learning platform*. Available at: <https://moodle.org/> (Accessed: 19 September 2021).

Moore, J.L., Dickson-Deane, C. and Galyen, K. (2010) 'E-Learning, online learning, and distance learning environments: Are they the same?', *Internet and Higher Education* [Preprint]. doi:10.1016/j.iheduc.2010.10.001.

Moseley, D., Baumfield, V., Elliot, J., Gregson, M., Higgins, S., Miller, J. and Newton, Douglas,

P. (2005) *Frameworks for Thinking*. Cambridge: Cambridge University Press.

Mulla, M.Z. (2020) *Cost, Activation, Loss Function: Neural Network: Deep Learning. What are these?*, Medium. Available at: <https://medium.com/@zeeshanmulla/cost-activation-loss-function-neural-network-deep-learning-what-are-these-91167825a4de> (Accessed: 12 November 2021).

Nargesian, F., Samulowitz, H., Khurana, U., Khalil, E.B. and Turaga, D. (2017) 'Learning feature engineering for classification', *IJCAI International Joint Conference on Artificial Intelligence*, 0(November), pp. 2529–2535. doi:10.24963/ijcai.2017/352.

Nassih, R. and Berrado, A. (2020) 'State of the art of Fairness, Interpretability and Explainability in Machine Learning: Case of PRIM', in *13th International Conference on Intelligent Systems: Theories and Applications*. Rabat, Morocco: ACM. doi:10.1145/3419604.3419776.

Newborn, M. (2000) 'Deep Blue's contribution to AI', *Annals of Mathematics and Artificial Intelligence*, 28, pp. 27–30.

Newcastle University (2015a) *Case Study: Commercial / Numbas*. Available at: <https://www.numbas.org.uk/sector/commercial/> (Accessed: 5 November 2020).

Newcastle University (2015b) *Case Study: Schools / Numbas*. Available at: <https://www.numbas.org.uk/sector/schools/> (Accessed: 5 November 2020).

Newcastle University (2015c) *Numbas*. Available at: <https://www.numbas.org.uk/> (Accessed: 2 November 2020).

Newcastle University (2018) *Case Studies, Numbas*. Available at: <https://www.numbas.org.uk/case-studies/> (Accessed: 16 September 2021).

Newcastle University (2021) *Form Builder, NUIT*. Available at: <https://services.ncl.ac.uk/itservice/services/collaboration/formbuilder/> (Accessed: 29

November 2021).

Newton, P.M. and Miah, M. (2017) 'Evidence-based higher education - Is the learning styles "myth" important?', *Frontiers in Psychology*, 8(MAR), pp. 1–9. doi:10.3389/fpsyg.2017.00444.

Nicol, D. and MacFarlane-Dick, D. (2006) 'Formative assessment and selfregulated learning: A model and seven principles of good feedback practice', *Studies in Higher Education*, 31(2), pp. 199–218. doi:10.1080/03075070600572090.

Nielsen, J. (1994) *10 Usability Heuristics for User Interface Design*, Nielsen Norman Group. Available at: <https://www.nngroup.com/articles/ten-usability-heuristics/> (Accessed: 18 November 2020).

Nielsen, M. (2015) 'How the Backpropagation Algorithm Works', in *Neural Networks and Deep Learning*. Determination Press.

Nilson, L. (2013a) *Creating Self-Regulated Learners: Strategies To Strengthen Students' Self-Awareness and Learning Skills*. Virginia: Stylus Publishing.

Nilson, L. (2013b) 'Fostering Self-Regulated Learning from the Start', in *Creating Self-Regulated Learners*. 1st Ed. Virginia: Stylus Publishing.

Noble, W.S. (2006) 'What is a Support Vector Machine?', *Nature Biotechnology*, 24(12), pp. 1565–1567. doi:10.1038/nbt1206-1565.

Northumbria University (2021) *Study at Northumbria*. Available at: <https://www.northumbria.ac.uk/study-at-northumbria/> (Accessed: 23 November 2021).

Nuhfer, E. and Knipp, D. (2003) 'The Knowledge Survey: A Tool For All Reasons', *To Improve The Academy*, 21(11), pp. 59–78. doi:0.1002/j.2334-4822.2003.tb00381.x.

NVIDIA (2020) *Deep Learning / NVIDIA Developer*. Available at: <https://developer.nvidia.com/deep-learning> (Accessed: 7 August 2020).

O'Grady, M.J. (2012) 'Practical problem-based learning in computing education', *ACM Transactions on Computing Education*, 12(3), pp. 1–16. doi:10.1145/2275597.2275599.

O'Leary, D.E. (2016) 'Ethics for Big Data and Analytics', *IEEE Intelligent Systems*, 31(4), pp. 81–84. doi:10.1109/MIS.2016.70.

Office for Students (2021a) *New analysis shows big boost in numbers of women enrolling on AI and data science courses*. Available at: <https://www.officeforstudents.org.uk/news-blog-and-events/press-and-media/new-analysis-shows-big-boost-in-numbers-of-women-enrolling-on-ai-and-data-science-courses/> (Accessed: 14 July 2021).

Office for Students (2021b) *The regulatory framework for higher education in England*. Available at: <https://www.officeforstudents.org.uk/advice-and-guidance/regulation/the-regulatory-framework-for-higher-education-in-england/> (Accessed: 15 July 2021).

Oliver, R. (2001) 'Developing e-learning environments that support knowledge construction in higher education', *Working for excellence in the e-economy.*, 1(1), pp. 407–416. Available at: <http://ro.ecu.edu.au/cgi/viewcontent.cgi?article=5765&context=ecuworks&sei-redir=1&referer=http%3A%2F%2Fscholar.google.com%2Fscholar%3Fhl%3Den%26q%3DDev%2Belo%2Be%2Blearning%2Benvironments%2Bthat%2Bsupport%2Bknowledge%2Bconstruction%2Bin%2Bhigh%2Beducation>.

Padlet (2021) *Padlet (about)*. Available at: <https://padlet.com/about> (Accessed: 21 September 2021).

Pallant, J. (2016) *SPSS Survival Manual*. 6th Editio. Berkshire: Open University Press.

Parker, P.D., Marsh, H.W., Ciarrochi, J., Marshall, S. and Abduljabbar, A.S. (2014) 'Juxtaposing math self-efficacy and self-concept as predictors of long-term achievement outcomes', *Educational Psychology*, 34(1), pp. 29–48. doi:10.1080/01443410.2013.797339.

Parmar, R. (2018) *Common Loss functions in Machine Learning, Towards Data Science*.

Available at: <https://towardsdatascience.com/common-loss-functions-in-machine-learning-46af0ff4d23> (Accessed: 12 November 2021).

Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems*. San Francisco: Morgan Kauffman Publishers.

Penprase, B.E. (2018) 'The Fourth Industrial Revolution and Higher Education', in Gleason. W, N. (ed.) *Higher Education In the Era Of The Fourth Industrial Revolution*. Singapore: Palgrave Macmillan.

Perfect, C. (2015) 'A demonstration of Numbas, an assessment system for mathematical disciplines', *CAA Conference* [Preprint].

Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz Harvard, B., Lyons, T., Manyika, J., Carlos Niebles, J. and Mishra, S. (2019) *Artificial Intelligence Index 2019 Annual Report, AI Index Steering Committee, Standford University - Human-Centered Artificial Intelligence*. Available at: https://hai.stanford.edu/sites/g/files/sbiybj10986/f/ai_index_2019_report.pdf.

Peters, M.A. (2017) 'Technological unemployment: Educating for the fourth industrial revolution', *Educational Philosophy and Theory*, 49(1), pp. 1–6. doi:10.1080/00131857.2016.1177412.

Peters, M.A. et al. (2020) 'Reimagining the new pedagogical possibilities for universities post-Covid-19', *Educational Philosophy and Theory*, 0(0), pp. 1–44. doi:10.1080/00131857.2020.1777655.

Petry, B., Mouton, H. and Reigeluth, C.M. (1987) 'A Lesson Based on the Gagné-Briggs Theory of Instruction', in Reigeluth, C.. (ed.) *Instructional Theories in Action*. New Jersey: Lawrence Erlbaum Associates.

Poon, J. (2013) 'Multimedia Education Resource for Learning and Online Teching (MERLOT) Blended Learning: An Institutional Approach for Enhancing Students' Learning Experiences',

Journal of Online Learning and Teaching, 9(2), pp. 271–288. Available at: <http://hdl.handle.net/10536/DRO/DU:30057995>.

Prensky, M. (2001) 'Digital Natives, Digital Immigrants', *On the Horizon*, 9(5), pp. 1–6.

Python Software Foundation (2021) *Python*, [Python.org](https://www.python.org/). Available at: <https://www.python.org/> (Accessed: 13 December 2021).

QAA (2019) *Subject Benchmark Statement: Computing*. Available at: <http://www.qaa.ac.uk/Publications/InformationAndGuidance/Documents/computing07.pdf>.

Qarabash, H.A.K. (2018) *Tools and Techniques for Exploring Engagement in a Flipped Classroom*. Newcastle University.

Quinlan, J.R. (1986) 'Induction of Decision Trees', *Machine Learning*, 1(1), pp. 81–106. doi:10.1007/bf00116251.

Race, P. (2020) *The Lecturer's Toolkit: A Practical Guide to Assessment, Learning and Teaching*. Fifth. Oxon: Routledge.

Rainford, J. (2020) 'Working with/in institutions: how policy enactment in widening participation is shaped through practitioners' experience', *British Journal of Sociology of Education*, 42(2), pp. 287–303. doi:10.1080/01425692.2020.1865130.

Ramalingam, V., LaBelle, D. and Wiedenbeck, S. (2004) 'Self-efficacy and Mental Models in Learning to Program', *Proceedings of the 9th Annual SIGCSE Conference on Innovation and Technology in Computer Science Education*, pp. 171–175. doi:10.1145/1007996.1008042.

Ramirez, G., Shaw, S.T. and Maloney, E.A. (2018) 'Math Anxiety: Past Research, Promising Interventions, and a New Interpretation Framework', *Educational Psychologist*, 53(3), pp. 145–164. doi:10.1080/00461520.2018.1447384.

Reinking, D., Labbo, L. and McKenna, M. (2000) 'From assimilation to accommodation: A

developmental framework for integrating digital technologies into literacy research and instruction', *Journal of Research in Reading*, 23(2), pp. 110–122. doi:10.1111/1467-9817.00108.

ResearchGate (2021) *Laura Heels's research works*, ResearchGate. Available at: <https://www.researchgate.net/scientific-contributions/Laura-Heels-2152010315> (Accessed: 20 September 2021).

Richards, N. and King, J. (2014) *Big Data Ethics*, *Wake Forest Law Review*. doi:10.1038/132817a0.

Richardson, F.C. and Suinn, R.M. (1972) 'The Mathematics Anxiety Rating Scale: Psychometric data', *Journal of Counseling Psychology*, 19(6), pp. 551–554. doi:10.1037/h0033456.

Rish, I. (2001) 'An Empirical Study of the Naive Bayes Classifier', *IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence*, 3(22).

Ritchie, J., Lewis, J., McNaughton Nicholls, C. and Ormston, R. (2014) *Qualitative Research Methods in Education*. 2nd edn. London: SAGE Publications.

Robins, A., Rountree, J. and Rountree, N. (2003) 'Learning and Teaching Programming: A Review and Discussion', *Computer Science Education*, 13(2), pp. 137–172. doi:10.1076/csed.13.2.137.14200.

Rumelhart, D., Hinton, G. and Williams, R. (1986) 'Learning Representations by Back-Propagating Errors', *Nature*, 323, pp. 533–536.

Rumelhart, D.E. and McClelland, J.L. (1987) *Learning Internal Representations by Error Propagation, Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations*. Massachusetts: MIT Press. doi:10.7551/mitpress/5236.001.0001.

Russell Group (2021) *Our Universities*. Available at: <https://russellgroup.ac.uk/about/our-universities/> (Accessed: 21 November 2021).

Russell, S. (2019) *Human Compatible: Artificial Intelligence And The Problem Of Control*. New York: Viking.

Russell, S. and Norvig, P. (2013) *Artificial Intelligence: A Modern Approach*. 3rd edn. England: Pearson.

Salkind, N. (2010) *Encyclopedia of Research Design - Triangulation*. California. doi:10.4135/9781412961288.

Saltz, J., Skirpan, M., Fiesler, C., Gorelick, M., Yeh, T., Heckman, R., Dewar, N. and Beard, N. (2019) 'Integrating ethics within machine learning courses', *ACM Transactions on Computing Education*, 19(4). doi:10.1145/3341164.

Sarkar, A., Dowker, A. and Kadosh, R.C. (2014) 'Cognitive enhancement or cognitive cost: Trait-specific outcomes of brain stimulation in the case of mathematics anxiety', *Journal of Neuroscience*, 34(50), pp. 16605–16610. doi:10.1523/JNEUROSCI.3129-14.2014.

Saygin, A.P., Cicekli, I. and Akman, V. (2000) 'Turing Test: 50 Years Later', *Minds and Machines*, 10, pp. 463–518. doi:10.1007/978-94-010-0105-2_2.

Scharff, L., Draeger, J., Verpoorten, D., Devlin, M., Dvorakova, S., Lodge, M., J. and Smith, S. (2017) 'Exploring Metacognition as Support for Learning Transfer', *Teaching and Learning Enquiry*, 5(1), pp. 1–14.

Schraw, G. and Dennison, R. (1994) 'Assessing Metacognitive awareness', *Contemporary Educational Psychology*, 19, pp. 460–475.

Schwab-McCoy, A., Baker, C.M. and Gasper, R.E. (2021) 'Data Science in 2020: Computing, Curricula, and Challenges for the Next 10 Years', *Journal of Statistics and Data Science Education*, 29(S1), pp. S40–S50. doi:10.1080/10691898.2020.1851159.

Schwab, K. (2017) *The Fourth Industrial Revolution*. London: Portfolio Penguin.

Scott, P. (2012) *It's 20 years since polytechnics became universities – and there's no going back*, *The Guardian*. Available at: <https://www.theguardian.com/education/2012/sep/03/polytechnics-became-universities-1992-differentiation> (Accessed: 7 December 2020).

Selby, C., Dorling, M. and Woollard, J. (2015) 'Evidence of assessing computational thinking', *IFIP TC3 Working Conference: a New Culture of Learning: Computing and Next Generations*, pp. 232–242. Available at: http://www.ifip2015.mii.vu.lt/file/repository/IFIP_Proceedings.pdf.

Select Committee on Artificial Intelligence (2018) *AI in the UK: ready, willing and able?*, *Aurthority of the House Of Lords*. doi:10.1177/0825859788004001-209.

Serbanescu, R. (2017) 'Identifying Threshold Concepts in Physics: too many to count!', *Practice and Evidence of the Scholarship of Teaching and Learning in Higher Education*, 12(2), pp. 378–396. Available at: <http://community.dur.ac.uk/pestlhe.learning/index.php/pestlhe/article/view/178/207>.

Shapiro, R.B., Fiebrink, R. and Norvig, P. (2018) 'How Machine Learning Impacts the Undergraduate Computing Curriculum', *Communications of the ACM*, 61(11), pp. 27–29. doi:10.1145/3277567.

Sheard, J. and Dick, M. (2011) 'Computing student practices of cheating and plagiarism: A decade of change', *ITiCSE'11 - Proceedings of the 16th Annual Conference on Innovation and Technology in Computer Science*, pp. 233–237. doi:10.1145/1999747.1999813.

Shinners-Kennedy, D. and Fincher, S.A. (2013) 'Identifying Threshold Concepts: From Dead End to a New Direction', *Proceedings of the ninth annual international ACM conference on International computing education research - ICER '13*, p. 9. doi:10.1145/2493394.2493396.

Shulman, L.S. (2013) 'Those who Understand: Knowledge Growth in Teaching', *Journal of Education*, 193(3), pp. 1–11. doi:10.1177/002205741319300302.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifra, L., van den Driessche, G., Graepel, T. and Hassabis, D. (2017) 'Mastering the game of Go without human knowledge', *Nature*, 550. doi:<https://doi.org/10.1038/nature24270>.

Silver, D. and Hassabis, D. (2017) *AlphaGo Zero: Starting from scratch*, DeepMind. Available at: <https://deepmind.com/blog/article/alphago-zero-starting-scratch> (Accessed: 23 July 2020).

Smith, T.W. and Colby, S.A. (2007) 'Teaching for Deep Learning', *The Clearing House: A Journal of Educational Strategies, Issues and Ideas*, 80(5), pp. 205–210. doi:[10.3200/tchs.80.5.205-210](https://doi.org/10.3200/tchs.80.5.205-210).

Sotola, L.K. and Crede, M. (2021) 'Regarding Class Quizzes: a Meta-analytic Synthesis of Studies on the Relationship Between Frequent Low-Stakes Testing and Class Performance', *Educational Psychology Review*, 33(2), pp. 407–426. doi:[10.1007/s10648-020-09563-9](https://doi.org/10.1007/s10648-020-09563-9).

Sprent, P. (2014) 'Fisher Exact Test', *International Encyclopedia of Statistical Science*, pp. 524–525. doi:[10.1007/978-3-642-04898-2_253](https://doi.org/10.1007/978-3-642-04898-2_253).

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. and Salakhutdinov, R. (2014) 'Dropout: A Simple Way to Prevent Neural Networks from Overfitting', *Journal of Machine Learning Research*, 15. doi:[10.1016/0370-2693\(93\)90272-J](https://doi.org/10.1016/0370-2693(93)90272-J).

Stack Exchange Inc (2022) *Stack Overflow*. Available at: <https://stackoverflow.com/> (Accessed: 17 January 2022).

Stanford University (2021) *Artificial Intelligence Index Report*. doi:[10.1145/3399971.3399984](https://doi.org/10.1145/3399971.3399984).

Stanford University (2022) *Stanford University*. Available at: <https://www.stanford.edu/> (Accessed: 15 January 2022).

Stankov, L., Lee, J., Luo, W. and Hogan, D.J. (2012) 'Confidence: A better predictor of academic achievement than self-efficacy, self-concept and anxiety?', *Learning and Individual*

Differences, 22(6), pp. 747–758. doi:10.1016/j.lindif.2012.05.013.

Stankov, L., Morony, S. and Lee, Y.P. (2014) 'Confidence: The best non-cognitive predictor of academic achievement?', *Educational Psychology*, 34(1), pp. 9–28. doi:10.1080/01443410.2013.814194.

Stanton, J.M. (2001) 'Galton, Pearson, and the Peas: A Brief History of Linear Regression for Statistics Instructors', *Journal of Statistics Education*, 9(3). doi:10.1080/10691898.2001.11910537.

Stathoulopoulos, K. and Mateos-Garcia, J.C. (2019) *Gender Diversity in AI Research*, Nesta. doi:10.2139/ssrn.3428240.

Stead, D.R. (2005) 'A review of the one-minute paper', *Active Learning in Higher Education*, 6(2), pp. 118–131. doi:10.1177/1469787405054237.

Steele, J., Holbeck, R. and Mandernach, J. (2019) 'Defining Effective Online Pedagogy', *Journal of Instructional Research*, 8(2), pp. 5–8. doi:10.9743/jir.2019.8.2.1.

STEM Women (2021) *Women in STEM | Percentages of Women in STEM Statistics*, STEM Women. Available at: <https://www.stemwomen.co.uk/blog/2021/01/women-in-stem-percentages-of-women-in-stem-statistics> (Accessed: 1 April 2021).

Sterrett, S.G. (2000) 'Turing's two tests for intelligence', *Minds and Machines*, 10(4), pp. 541–559. doi:10.1023/A:1011242120015.

Steup, M. and Neta, R. (2020) *Epistemology*, Stanford Encyclopedia of Philosophy. Available at: <https://plato.stanford.edu/entries/epistemology/> (Accessed: 20 July 2021).

Steven, K. and Thomas, L. (2019) *Attracting Diversity : End of Project Report*, Advance HE. Available at: <https://www.advance-he.ac.uk/knowledge-hub/attracting-diversity-end-project-report> (Accessed: 3 January 2020).

Stierer, B. and Antoniou, M. (2004) 'Are there distinctive methodologies for pedagogic research in higher education?', *Teaching in Higher Education*, 9(3), pp. 275–285. doi:10.1080/1356251042000216606.

Storey, M.A.D., Fracchia, F.D. and Müller, H.A. (1999) 'Cognitive Design Elements to Support the Construction of a Mental Model During Software Exploration', *Journal of Systems and Software*, 44(3), pp. 171–185. doi:10.1016/S0164-1212(98)10055-9.

Sutherland, P. (1999) 'The application of Piagetian and Neo-Piagetian ideas to further and higher education', *International Journal of Lifelong Education*, 18(4), pp. 286–294. doi:10.1080/026013799293702.

Tabachnick, B.G. and Fidell, L.S. (2013) *Using Multivariate Statistics*. 6th edn. Boston: Pearson Education.

Taylor, J.A. and Mohr, J. (2001) 'Mathematics for Math Anxious Students Studying at a Distance', *Journal of Developmental Education*, 25(1), pp. 30–37. Available at: <http://web.a.ebscohost.com/sas/detail?sid=c1259a70-60b0-4cf7-a17e-b9cf790fc1cf%40sessionmgr4005&vid=0&hid=4214&bdata=JkF1dGhUeXBIPQ%3D%3D#AN=5223062&db=afh>.

Tendre, M. and Apiola, M. (2013) 'Three Computing Traditions In School Computing Education', in Kadijevich, D., Angeli, C., and Schulte, C. (eds) *Improving Computer Science Education*. New York: Routledge.

The Guardian (2020) *The Best UK Universities 2021 – Rankings*, *The Guardian*. Available at: <https://www.theguardian.com/education/ng-interactive/2020/sep/05/the-best-uk-universities-2021-league-table> (Accessed: 5 December 2020).

The Observer (2003) *Ivy League for the UK*. Available at: <https://www.theguardian.com/education/2003/sep/21/highereducation.uk1> (Accessed: 2 July 2020).

The Royal Society (2017) *Machine Learning: The Power and Promise of Computers that Learn by Example, Report by the Royal Society*. doi:10.1126/scitranslmed.3002564.

The University of Manchester (2021) *The University of Manchester*. Available at: <https://www.manchester.ac.uk/> (Accessed: 13 December 2021).

Theodoridis, S. (2020) 'Classification: a Tour of the Classics', *Machine Learning*, pp. 301–350. doi:10.1016/B978-0-12-818803-3.00016-7.

Thomas, D.R. (2006) 'A General Inductive Approach for Analyzing Qualitative Evaluation Data', *American Journal of Evaluation*, 27(2), pp. 237–246. doi:10.1177/1098214005283748.

Times Higher Education (2021) *World University Rankings 2021 by subject: computer science*, Times Higher Education. Available at: https://www.timeshighereducation.com/world-university-rankings/2021/subject-ranking/computer-science#!/page/0/length/25/sort_by/rank/sort_order/asc/cols/stats (Accessed: 11 February 2021).

Tripathi, M. (2020) *Underfitting and Overfitting in Machine Learning*, Data Science Foundation. Available at: <https://datascience.foundation/sciencewhitepaper/underfitting-and-overfitting-in-machine-learning> (Accessed: 11 November 2021).

Turing, A. (1950) 'Computing Machinery and Intelligence', *Mind*, LIX(236), pp. 433–460.

Udacity (2020) *The School Of Artificial Intelligence*. Available at: <https://www.udacity.com/school-of-ai> (Accessed: 15 October 2020).

Udacity (2021) *The Latest Digital Skills Within Reach*. Available at: <https://www.udacity.com/>.

Udacity (2022) *AI for Beginners*. Available at: <https://www.udacity.com/course/artificial-intelligence--ud954> (Accessed: 22 May 2022).

UK AI Council (2021) *AI Roadmap*, GOV.UK. Available at:

<https://www.gov.uk/government/publications/ai-roadmap> (Accessed: 1 March 2021).

United Kingdom and Ireland Computing Education Research (2021) *UKICER 2021, ACM*. Available at: <https://www.ukicer.com/> (Accessed: 20 September 2021).

University College London (2021) *Machine Learning MSc*. Available at: <https://www.ucl.ac.uk/prospective-students/graduate/taught-degrees/machine-learning-msc> (Accessed: 9 November 2021).

University of Bath (2021) *University of Bath*. Available at: <https://www.bath.ac.uk/> (Accessed: 23 November 2021).

University of Cambridge (2020a) *Department of Computer Science and Technology – Course pages 2020–21: Data Science: principles and practice*. Available at: <https://www.cl.cam.ac.uk/teaching/2021/DataSciII/> (Accessed: 9 November 2021).

University of Cambridge (2020b) *Principles of Machine Learning Systems*. Available at: <https://www.cl.cam.ac.uk/teaching/2021/L46/> (Accessed: 25 November 2021).

University of Edinburgh (2018) *Course Catalogue - Machine Learning Practical*. Available at: <http://www.drps.ed.ac.uk/17-18/dpt/cxinfr11132.htm> (Accessed: 9 November 2021).

University of Manchester (2021) *Undergraduate Syllabus (School of Computer Science)*. Available at: <https://studentnet.cs.manchester.ac.uk/ugt/COMP24112/syllabus/> (Accessed: 9 November 2021).

University of St Andrews (2021a) *CS5011 - Artificial Intelligence Practice* University of St Andrews. Available at: https://www.st-andrews.ac.uk/subjects/modules/catalogue/?code=CS5011&academic_year=2021/2 (Accessed: 9 November 2021).

University of St Andrews (2021b) *University of St Andrews - Scotland's first university, founded 1413*. Available at: <https://www.st-andrews.ac.uk/> (Accessed: 14 December 2021).

University of Sunderland (2021) *MSc Data Science*. Available at: <https://www.sunderland.ac.uk/study/computing/postgraduate-data-science/#course-structure> (Accessed: 9 November 2021).

University of Waikato (2020) *WEKA- The workbench for machine learning*. Available at: <https://www.cs.waikato.ac.nz/ml/weka/> (Accessed: 26 November 2020).

Vartiainen, H., Tedre, M. and Valtonen, T. (2020) 'Learning machine learning with very young children: Who is teaching whom?', *International Journal of Child-Computer Interaction*, 25, p. 100182. doi:10.1016/j.ijCCI.2020.100182.

Volet, S. (1991) 'Modelling and coaching of relevant metacognitive strategies for enhancing university students' learning', *Learning and Instruction*, 1(4), pp. 319–336.

Walker, G. (2013) 'A cognitive approach to threshold concepts', *Higher Education*, 65(2), pp. 247–263. doi:10.1007/s10734-012-9541-4.

Watermeyer, R., Crick, T., Knight, C. and Goodall, J. (2020) 'COVID-19 and digital disruption in UK universities: afflictions and affordances of emergency online migration', *Higher Education* [Preprint]. doi:10.1007/s10734-020-00561-y.

Watkins, C. and Dayan, P. (1992) 'Q-Learning', *Machine Learning*, 8, pp. 279–292.

Wenger, E. (2018) 'A Social Theory of Learning', in Illeris, K. (ed.) *Contemporary Theories of Learning*. New York: Routledge.

Werner, K.M. and Milyavskaya, M. (2019) 'Motivation and self-regulation: The role of want-to motivation in the processes underlying self-regulation and self-control', *Social and Personality Psychology Compass*, 13(1), pp. 1–14. doi:10.1111/spc3.12425.

World Economic Forum (2016) *The Future of Jobs: Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution*. doi:10.23943/princeton/9780691172811.003.0009.

World Economic Forum (2018) *The Global Gender Gap Report 2018*. Available at:

<https://www.weforum.org/reports/the-global-gender-gap-report-2018> (Accessed: 14 July 2021).

World Health Organization (2021) *Coronavirus disease (COVID-19)*. Available at: <https://www.who.int/emergencies/diseases/novel-coronavirus-2019> (Accessed: 4 August 2021).

Worster, A., Fan, J. and Ismaila, A. (2007) 'Understanding Linear and Logistic Regression Analyses', *Journal of Emergency Medicine*, 9(2), pp. 111–113.

Yang, P. and Cheng, Y. (2018) 'Educational Mobility and Transnationalization', in Gleason. W, N. (ed.) *Higher Education In the Era Of The Fourth Industrial Revolution*. Singapore: Palgrave Macmillan.

Yang, X., Zhang, M., Kong, L., Wang, Q. and Hong, J.C. (2021) 'The Effects of Scientific Self-efficacy and Cognitive Anxiety on Science Engagement with the "Question-Observation-Doing-Explanation" Model during School Disruption in COVID-19 Pandemic', *Journal of Science Education and Technology*, 30(3), pp. 380–393. doi:10.1007/s10956-020-09877-x.

Yeo, A., Legard, R., Keegan, J., Ward, K., McNaughton Nicholls, C. and Lewis, J. (2014) 'In-Depth Interviews', in Ritchie, J. et al. (eds) *Qualitative Research Practice*. London: SAGE Publications.

Young, E., Wajcman, J. and Sprejer, L. (2021) *Where are the women ? Mapping the Gender Job Gap in AI.*, Alan Turing Institute. Available at: https://www.turing.ac.uk/research/publications/report-where-are-women-mapping-gender-job-gap-ai?_cldee=Yi5hbGxlbjJAbmV3Y2FzdGxlLmFjLnVr&recipientid=contact-37781e7a802be811811a70106faad2f1-b6a1772674d34b249cfb106c2a7b42d5&esid=e6bc640a-8090-eb11-b1ac-000d (Accessed: 5 June 2021).

Zaharias, P. and Poylymenakou, A. (2009) 'Developing a usability evaluation method for e-learning applications: Beyond functional usability', *International Journal of Human-Computer Interaction*, 25(1), pp. 75–98. doi:10.1080/10447310802546716.

Zar, J.H. (2005) 'Spearman Rank Correlation', *Encyclopedia of Biostatistics* [Preprint]. doi:10.1002/0470011815.B2A15150.

Zepke, N. (2013) 'Threshold concepts and student engagement: Revisiting pedagogical content knowledge', *Active Learning in Higher Education*, 14(2), pp. 97–107. doi:10.1177/1469787413481127.

Zhang, W.J., Yang, G., Lin, Y., Ji, C. and Gupta, M.M. (2018) 'On Definition of Deep Learning', *World Automation Congress Proceedings*, 2018-June, pp. 232–236. doi:10.23919/WAC.2018.8430387.

Zheng, A. and Casari, A. (2018) *Feature Engineering for Machine Learning*. California: O'Reilly.

Zoom Video Communications Inc (2021) *Zoom*. Available at: <https://zoom.us/> (Accessed: 11 November 2021).

Appendix A: Online Questionnaire for Lecturers

The following link is to the questionnaire sent to lecturers who teach some form of AI course: https://data.ncl.ac.uk/articles/workflow/Lecturer_Questionnaire/16587020/1

Appendix B: Interview Questions for Lecturers

1. At what level is Machine Learning taught?
2. How is the module delivered? (Block taught, weekly, practical sessions)
3. Any course prerequisites?
4. What level of maths skills do you expect the students to have?
5. What level of programming knowledge do you expect the students to have?
6. What Machine Learning topics are covered within the module?
7. Is there any particular aspect of the module that you focus a lot of time on as you think it is a pivotal topic?
8. Do you teach any Deep Learning?
9. How is the module assessed?
10. Is there any particular aspect that you feel students struggle with?
11. Where would you advise students to look for additional information?
12. Do you receive student feedback on the module?
13. Module outcome – pass/fail percentage
14. Do you enjoy teaching the module?
15. Any specific teaching strategies you employ?

Appendix C: Pre-Module Student Questionnaire

Machine Learning Questionnaire

1. Background Information

a) Age:

Please tick one box

18 -22 years old 23-27 years old 28-32 years old
33-40 years old 40+ years old Prefer not to say

b) Gender:

Male Female Prefer not to say

I prefer to self-describe: _____

c) Please state your previous degree: _____

d) Did you study this degree at Newcastle University?

Yes No

e) How would you describe your level of maths attainment?

Studied maths up to GCSE level or equivalent

Studied maths up to A-Level or equivalent

Maths was a major part of my first degree

My first degree was in maths

f) On a scale of 1-10 how confident would you describe yourself within your maths knowledge?

Please circle a number

1	2	3	4	5	6	7	8	9	10
Not confident					Very confident				

g) How would you describe your level of programming skills?

Beginner Novice Expert

h) On a scale of 1-10 how confident would you describe yourself in the application of your programming skills?

1	2	3	4	5	6	7	8	9	10
Not confident					Very confident				

2. Machine Learning

a) Please write a brief description of what you think machine learning is: _____

b) On a scale of 1-10 how interested are you in studying machine learning?

1	2	3	4	5	6	7	8	9	10
Not very interested									Very interested

c) On a scale of 1-10 how confident do you feel in your ability to do well in this module?

1	2	3	4	5	6	7	8	9	10
Not confident									Very confident

d) Which aspect of the module do you expect to find the most difficult?

Theoretical subject knowledge / exam

Practical exercises / coursework

3. Teaching

a) Are there any skills from previous learning which you think will aid within the machine learning module?

Please tick all that apply

Team learning <input type="checkbox"/>	Memorisation <input type="checkbox"/>	Motivation <input type="checkbox"/>	Resilience <input type="checkbox"/>
Organisation <input type="checkbox"/>	Perseverance <input type="checkbox"/>	Programming skills <input type="checkbox"/>	
Numeracy skills <input type="checkbox"/>	Research skills <input type="checkbox"/>	Analytical skills <input type="checkbox"/>	

b) How important would you rate these resources in aiding your learning?

Please rate these from 1-5, with 1 being the most important

Resource	Rating (1-5)
Lectures	
Handouts / written material	
Quizzes	
Practical exercises	
Assignments	

c) Please state any additional resources which you think would aid your learning:

d) In order of priority, rate where you would most likely go for additional support:

Please rate these from 1-6, with 1 being the first choice

Support Option	Rating (1-6)
Module leader	
Demonstrator	
Peers	
Textbook	
Online	

Other (please state) : _____	_____
------------------------------	-------

e) Which of these learning strategies do you propose to use in this module:

Please tick all which apply

Note taking Study group Practical exercises Quizzes

Textbooks online guidance Critical thinking Reflection

Goal setting Planning Self-evaluation

Other (Please state): _____

4. Research

As part of this research I am looking to interview students and hold focus groups regarding their experience in learning machine learning, if you would like to participate or require further information on this study please contact B.Allen2@newcastle.ac.uk

Appendix D: Observation Guide

1. Background

Location: _____

Date: _____

Observation start time: _____ End time: _____

2. Lecture

Type of lecture	Mainly structured / Mainly unstructured / Combination
Purpose of lecture	
Topics planned to be covered	
Lecture context notes (attendance, set up, male/female ratio)	
Additional Comments	

3. Physical Environment

Any changes or specific details on layout / set up?	
---	--

4. Activity – Sequence of events and activities

Notes on what happens:	
------------------------	--

--	--

5. Any notes on interactions (simultaneous with 4 above)

Between staff and student	
Between peers	

6. Closing

What happens next?	
---------------------------	--

Task for practical?	

7. Post observation perceptions

<p><i>Did actual activity differ from any plans?</i> Researcher perception</p>	
--	--

8. Any other post-observation notes

- Possible lines of further enquiry
- Reflections on the methods used in observation
- Ethical issues, tensions, problems
- Observer reaction

- **Points of clarification**



Appendix E: Week 2 One Minute Paper

Topics Covered this week:

- Support Vector Machines
- Decision Tree and Random Forest
- K-Nearest Neighbour Classifier
- Clustering

Which topic(s) did you find the most difficult this week?

Which topic are you still unsure on?

Any specific questions?

How confident do you feel in what you have learnt this week?

Please circle a number

1	2	3	4	5	6	7	8	9	10
Not confident					Very confident				

Appendix F: Student Post-Module Questionnaire

The following link is to the questionnaire the students were asked to complete on completion of their Machine Learning module at university A:

https://data.ncl.ac.uk/articles/workflow/Student_Questionnaire/16587017/1

The following link is to the questionnaire the students were asked to complete on completion of their Artificial Intelligence module at university B:

<https://forms.ncl.ac.uk/view.php?id=3190269>

Appendix G: Miscellaneous Machine Learning Modules

Miscellaneous Machine Learning Modules				
Institution	Title	Educational Level	Pre-Requisites	Content
Cambridge University	Machine Learning and Real-World Data	Undergraduate	None Listed	Statistical classification, sequence analysis, social networks
Cambridge University	Principles of Machine Learning Systems	Postgraduate	None Listed	Distributed learning algorithms, Deep Learning compilers
University of Glasgow	Machine Learning & Artificial Intelligence for Data Scientists	Postgraduate	None Listed	Regression, classification, clustering, decision making problems, ethical issues
University of Edinburgh	Machine Learning and Pattern Recognition	Undergraduate	Algebra, vectors/matrices, calculus, probability, programming skills	Classification, regression, neural networks, clustering, dimensionality reduction
University of Edinburgh	Machine Learning Practical	Undergraduate	Previous modules in Machine Learning, basic maths and probability, programming experience	Optimisation and learning rules, neural networks, autoencoders, CNNs and RNNs, regularisation and normalisation
University College London	Supervised Learning	Postgraduate	Calculus, probability, linear algebra	Linear regression, SVM, neural networks
University College London	Applied Machine Learning	Postgraduate	Linear algebra, calculus, probability, programming skills	Classification, regression, clustering, optimization methods
University of Kent	Introduction to Intelligent Systems	Undergraduate	Object-Oriented Programming module	Supervised and unsupervised learning, reinforcement learning
University of Birmingham	Machine Learning and Intelligent Data Analysis	Both Levels	Some knowledge of mathematics	PCA, clustering, classification
University of Sunderland	Machine Learning and Data Analytics	Postgraduate	None Listed	Ethical aspects, principles of modelling

Miscellaneous Machine Learning Modules				
Institution	Title	Educational Level	Pre-Requisites	Content
Northumbria University	Intelligent Systems	Undergraduate	None Listed	Data and text mining, NLP, data visualisation
Northumbria University	Machine Learning and Computer Vision	Undergraduate	None Listed	Supervised and unsupervised learning, legal, ethical and social issues
Newcastle University	Predictive Analytics and Machine Learning	Undergraduate	None Listed	Supervised learning, clustering, deep learning

Appendix H: Assessment from AI Module

Please read **all** instructions and information carefully.

This assignment contributes **30%** to your final module mark and will assess the following learning outcomes:

Knowledge

- Knowledge of a wide range of AI techniques, which are being applied in industry or research, allowing them to choose and apply the correct AI techniques for the problems which arise.
- Awareness of current and new/future developments in the field of AI and its applications.

Skills

- Assess real-world problems and determine which AI approaches are suitable for their solutions
- Apply various AI models and techniques in the solutions of a range of problems, and characterise the expected performance of a model, and compare with other techniques.

Introduction

This assignment will involve the development of a portfolio of practical work, within which you demonstrate your ability to design, implement, and evaluate an intelligent prototype for a selected scenario. Your application will be developed in the python programming language and will make use of current developments in the field AI, for example through utilisation of relevant APIs and methodologies.

You have been issued with exercises associated with lab work throughout the module and these should be uploaded to your e-Portfolio. Any work, research or planning that you carry out should also be uploaded to your e-Portfolio.

You may choose **one** of the following projects:

Project Prototype 1: Pathfinding with Planning and search

Project Prototype 2: Machine learning to solve real world problem

You will spend tutorials doing work which could contribute to each of these project areas. You will fully develop and submit **one** project prototype in the python programming language during the professional practice weeks. This assessment covers the initial planning stage of the prototype.

Prototype Planning

To be completed during Professional Practice Week 1:

Introduction (around 500 words)

The introduction should include a short mission statement for your proposed prototype and should provide a short overview of the contents to-date of your e-Portfolio on Canvas.

Section 1: Prototype Identification and Planning (around 1,000 words)

Section 1.1 Literature Review on Prototype Identification (around 700 words)

This section should be a literature review of your solution identification and software development planning. **Suggested themes** are – the background of AI, in particular the specialism associated with your chosen prototype development, how similar solutions have been employed and the success or otherwise of these examples.

Section 1.2 Reflection on the Prototype Identification (around 300 words)

This section should be a reflection on your experience of your activity at the end of the prototype identification and planning stage, drawing both from what you learned from the research activity as embodied in the literature review and from what you have learned from practical activities that have been carried out during the module. This is about your experience of the process during this early stage of identifying the scope of your application prototype.

Suggested timeframe for this work: Professional Practice Week 1 (**14 hours**)

Marking Scheme

Introduction	8 marks
Literature Review on Prototype Identification	14 marks
Reflection on the Prototype Identification	8 marks
Total	30 marks

Marking Criteria

Introduction	You outline your choice of prototype and a basic overview of your e-Portfolio work.	You discuss your chosen prototype and link it well to your e-Portfolio work.	You justify your choice of prototype using your e-Portfolio work as an evaluative argument for your choice	
0	1-3	4-6	7-8	
Literature Review	Basic description of your project choice, but lacking depth of research in the area.	Some discussion of project choice and some links to research in the area	Good discussion and justification of project choice with clear links to relevant research in the area.	Excellent argument, which justifies project choice with clear links to relevant research in the area.
0	1-4	5-7	8-11	12-14
Reflection	Basic reflection on your experience of the work undertaken	Good reflection on your experience, which links your assignment and practical tutorial work well	Strong reflection on both your assignment and practical work and how one has informed the other	
0	1-3	4-6	7-8	

Submission Instructions

Your document submission should be provided as a single document in either Word or PDF and uploaded to Canvas by the specified hand-in date, using the assignment submit link provided in the assessment area. You will only be able to submit the report to the **turn-it-in** portal **once**, so please make sure that you **only submit the final version**.

Your report should be accompanied by a reference list using the Harvard style of referencing and should use a good range of sources. To achieve a high mark, you will be expected to cite as least 4 academic references from conference proceedings or journals in your literature review. This will demonstrate that you have researched your chosen solution in depth.

Submission Date: Friday 13th November by 2pm via Canvas.

Appendix I: Links to MetaLearning Tutorials

Maths and Stats for Machine Learning:

<https://numbas.mathcentre.ac.uk/exam/10572/maths-and-stats-for-machine-learning/embed/?token=d633f234-aaa4-434c-b549-0c0eebf93a5>

Overview of AI:

<https://numbas.mathcentre.ac.uk/exam/10571/overview-of-ai/embed/?token=5a9e651c-b630-47be-a1d8-805a86145729>

Overview of Machine Learning:

<https://numbas.mathcentre.ac.uk/exam/10573/overview-of-machine-learning/embed/?token=3dafb449-5c98-4f43-a563-9e91b92f818a>

Machine Learning Algorithms:

<https://numbas.mathcentre.ac.uk/exam/11608/machine-learning-algorithms/embed/?token=afb0e720-2491-4fd0-a89d-580ebcebab23>

Data Preprocessing:

<https://numbas.mathcentre.ac.uk/exam/10886/data-preprocessing/embed/?token=bb06c1a7-bfca-4ace-a94f-e6858ed3bc43>

Overview of Deep Learning:

<https://numbas.mathcentre.ac.uk/exam/10575/overview-of-deep-learning/embed/?token=dbd012ec-1da4-48c1-bc3b-d10b12e57a70>

Appendix J: Student Questionnaire – MetaLearning Review

The following link is to the questionnaire to be completed by students relating to their experience using MetaLearning: <https://forms.ncl.ac.uk/view.php?id=6719176>

Appendix K: Maths and Stats Knowledge Survey

The following link is to the Maths and Stats knowledge survey which participants are asked to complete within the Maths and Stats tutorial: <https://forms.ncl.ac.uk/view.php?id=6710907>

Appendix L: Information Sheet for Lecturers Reviewing MetaLearning

Information Sheet

12/07/2021

Title of Study: Identifying the Best Practices for Teaching Machine Learning

Invitation and Brief Summary

You are being invited to take part in a research study. Before you decide whether or not you wish to take part it is important that you understand why the research is being done and what it will involve. Please read this information carefully and discuss it with others if you wish. Take time to decide whether or not you wish to take part. If you do decide to take part, you will be asked to sign a consent form supplied at the end of this information sheet. However, you are free to withdraw at any time, without giving any reason and without any penalty.

What is the purpose of the research?

It is already established that machine learning is a difficult topic to learn, however limited research has been undertaken which identifies the barriers that students face when learning this area of computer science and possible strategies to aid with the comprehension of this domain. As part of this study an online learning tool has been created to assist students in their learning of this domain. The tool includes 6 tutorials on:

- Maths and Stats for Machine Learning
- Data Preprocessing
- Overview of Artificial Intelligence
- Overview of Machine Learning
- Machine Learning Algorithms
- Overview of Deep Learning

Alongside the tutorials, strategies to help boost student self-efficacy and confidence in their skills are included such as low stakes quizzing and knowledge surveys. This research aims to identify the perceived usefulness of the mitigation strategies included and the use of an online tool to assist students learning this domain.

What does taking part involve?

The participant will be asked to view and briefly use the online machine learning tool, they will then attend a meeting where they will be interviewed regarding their experience and thoughts on the online machine learning tool. The interview will last approximately 30 minutes and the researcher will make notes throughout this time. The participant will be asked whether they consent to the interview being audio and potentially video recorded for transcription purposes. However, the recording of the interview is optional and will not impede on participation in the research.

What information will be collected and who will have access to the information collected?

The main information collected from the study will be your opinions on the online learning tool. We will use your name and email address to contact you about the research study. Individuals at Newcastle University may look at the research data to check the accuracy of the research study. The only individuals at Newcastle University who will have access to

information that identifies you will be individuals who need to contact you for follow up information such as research outcomes or to audit the data collection process.

Who is the sponsor and data controller for this research?

Newcastle University

Has this study received ethical approval?

This study received ethical approval from the Faculty of Science, Agriculture and Engineering Ethics Committee on 26/09/2018.

Who should I contact for further information relating to the research?

Principal Investigator: Becky Allen Email: B.Allen2@newcastle.ac.uk

PhD Supervisors: Dr Marie Devlin, Dr Stephen McGough

Please read the following instructions and confirm consent:

Please indicate with a Y/N (Yes/No) whether you confirm consent		
1.	I confirm that I have read the information above regarding the study and have had the opportunity to consider the information and ask questions.	
2.	I understand that my participation is voluntary and that I am free to withdraw from this study at any point without giving a reason. I understand that if I decide to withdraw, any data provided up to that point will be omitted.	
3.	I consent to the processing of my personal information [name, experience] for the purposes of this research study and understand that my data will be anonymised.	
4.	I understand that all information will be treated in confidence and will be disposed of on 10/01/2022.	
5.	I understand that my research data may be published but that my personal data and any identifying information will be anonymised.	
6.	I consent to be recorded for the interview and understand that the recordings will be stored anonymously on the researcher's computer and used for research purposes only.	
7.	I understand that recording of the interview is optional and therefore not necessary for my participation in the research.	
8.	I agree to take part in this research project.	

Name of participant:	Signature:	Date:

Appendix M: History of AI Timeline from MetaLearning

